Geometry Processing from an Elastic Perspective

Martin Rumpf, Max Wardetzky

February 27, 2014

Abstract

Triggered by development of new hardware for high resolution acquisition of complex geometric objects such as laser range scanners, new graphics processors for realtime rendering and animation of extremely detailed, textured geometric structures, and novel rapid prototyping equipment such as 3D printers the processing of highly resolved complex geometries has established itself as an important area both of fundamental research and impressive applications. Concepts from image processing have been picked up and carried over to surfaces, physical based modeling plays a crucial role, and aspects of computer aided geometry design have been integrated. This paper aims at highlighting some of the developments with a particular focus on methods related to the mechanics of thin elastic surfaces from the authors point of view. An overview of different geometric representations ranging from point clouds over level sets to subdivision surfaces is given together with a sketch of the basic computational tools. Furthermore, fundamental computational tasks such denoising, deformation or matching, and spectral methods are discussed. Finally, beyond the processing of single shapes, it will be described how spaces of shapes can be investigated using concepts from Riemannian geometry.

1 Introduction

Over the last decades the surface models used in computer graphics have been getting successively more complex. In particular, new hardware for the acquisition of geometries such as laser range scanners and new software for instance for multi view reconstruction provide high resolution geometric models with triangulations consisting of millions of triangles or point clouds consisting of an even larger set of point measurements. Depending on the origin of these models and the concrete purposes different geometric representations are appropriate to process, model or animate the underlying surface geometries. On this background the research field field of geometry processing has undergone a rapid development. In particular physical based modeling plays a crucial role and concepts from mechanics have been picked up and adapted to the needs in computer graphics.

The overarching goal of geometry processing is the combination of valid physical models with efficient, near realtime simulation. Recently, an established trend is to *not* sacrifice physical accuracy for computational efficiency. One important instance of this trend is to work with reliable nonlinear models and to use physical insight to accelerate algorithms instead of a brute force linearization. This provides a challenge for many applications. The diverse approaches that we are going to present and compare here all face this challenge in one way or the another. When asking for discrete geometric approximations of smooth surfaces, a central goal is to provide a structure preserving, consistent and converging notion of discrete curvatures and discrete elastic energies. For subdivision surfaces, an open problem is to link the design of subdivision algorithms with physical principles and to handle the high complexity of the underlying algorithms. If one votes for an implicit representation of surfaces via level set methods the computational workload raises from dimension two to three.

This paper highlights some of the ongoing developments with a particular focus on methods related to the mechanics of thin surfaces. As a disclaimer, let us emphasize that the exposition in this paper should not be read as an objective and balanced overview of geometry processing as a whole, but rather, as the title suggests, as a personal perspective strongly biased by an eye on continuum mechanics, variational methods, and partial differential equations. Indeed, we focus on surfaces which are primarily modeled as thin elastic shells and discuss the implementation of corresponding variational models. Furthermore, we make use of Rayleigh's paradigm by which viscous dissipation models are derived from elastic energies replacing elastic strains by strain rates. Even on this restricted research field we have to acknowledge that we are

unable to do any sort of justice to the plethora of ideas and developments. On the other hand, we discuss in parallel the different types of geometry representation and compare both the mathematical foundation and the algorithmic realization.

The paper is organized as follows. In Section 2 we review different representations of surfaces ranging from parametrized and triangulated surface over subdivision surfaces, level sets and point clouds to models from discrete geometry. In particular we show how to model and discretize the most basic geometric functionals and PDE operators such as the Dirichlet energy, the Laplacian and the Willmore functional. Thereby, this section lays the foundation for the exposition of the models and methods described in the subsequent sections. Then, in Section 3 we review some surface flows and surface fairing methods based on feature aware geometric diffusion. How to model the stored elastic energy of a surface using a thin shell approach is discussed in Section 6. We define membrane and bending energies on parametrized and implicit surfaces and discuss the approximation using discrete exterior calculus. Applications to surface matching and physical simulation are outlined. The spectral analysis of the Hessian of surface energies and the computation of eigenfrequences and vibration models are fundamental tools in surface modeling and will be briefly presented in Section 7. Furthermore, in Section 8 we anticipate a Riemannian perspective on the space of surfaces and sketch the set up of a time discrete Riemannian calculus. Thereby, the underlying metric represents the rate of physical dissipation accumulated along paths in shell space. Finally, we draw conclusions in Section 9.

2 Different geometry representations and PDE approaches

The focus of this article is on geometry processing tools based on partial differential equations and geometric functionals. In the application different tools are implemented on different types of geometry representation. To emphasize the similarities and relations of these approaches we will discuss in this section some of the most basic differential operators and energy functionals relevant in geometry processing, namely the Dirichlet energy, the Laplacian, and the Willmore functional. Furthermore, we will briefly sketch how they can be spatially discretized.

The flat case. We briefly recall the Euclidian case. On a two dimensional domain Ω the Dirichlet energy of a function $u : \Omega \to \mathbb{R}$ is defined as

$$\mathcal{W}_{\text{Dirichlet}}[u] = \frac{1}{2}\int_\Omega |\nabla u|^2 \,\mathrm{d}x\,.$$

The weak form of the (negative) Laplacian $-\triangle = -\mathrm{div}\nabla$ of a function u is given as the first variation of Dirichlet energy. Indeed, using integration by parts one easily verifies that

$$\int_\Omega -\triangle u\phi \,\mathrm{d}x = \partial_u \mathcal{W}_{\text{Dirichlet}}[u](\phi) = \int_\Omega \nabla u \cdot \nabla\phi \,\mathrm{d}x \tag{1}$$

for smooth u and for any smooth, compactly supported test function ϕ on Ω. A simple model for a thin plate energy on the plate domain Ω and for a vertical displacement u is given by $\mathcal{W}_{\text{plate}}[u] = \frac{1}{2}\int_\Omega |\triangle u|^2 \,\mathrm{d}x$. The first variation of this energy leads to the bi-Laplacian $\triangle^2 u$ with $\int_\Omega \triangle^2 u\phi \,\mathrm{d}x = \int_\Omega \triangle u \cdot \triangle\phi \,\mathrm{d}x$ for smooth u, or in weaker form

$$\int_\Omega \triangle^2 u\phi \,\mathrm{d}x = \int_\Omega \nabla w \cdot \nabla\phi \,\mathrm{d}x\,, \quad \int_\Omega w\phi \,\mathrm{d}x = \int_\Omega \nabla u \cdot \nabla\phi \,\mathrm{d}x\,, \tag{2}$$

with $\triangle^2 u = -\triangle w$ and $w = -\triangle u$ in the weak sense of (1).

Finite Element discretization in the flat case. With respect to a finite element discretization of the above energies and differential operators, one considers a regular, (for the ease of presentation) triangular mesh $\mathcal{T}_h$ covering the domain Ω, which we assume for simplicity to be polygonally bounded. Here h denotes the grid size defined as the maximal diameter of the triangles $T \in \mathcal{T}_h$. At first we consider the space $\mathcal{V}_h$ of piecewise affine, continuous functions U on Ω. Each function U is uniquely described by a vector $\bar{U}$ of nodal values on the vertices of the triangulation. Now, the variation of the Dirichlet energy $\mathcal{W}_{\text{Dirichlet}}$ on discrete functions U gives rise to a quadratic form $\mathbf{L}\bar{U}\cdot\bar{V} := \int_\Omega \nabla U \cdot \nabla V \,\mathrm{d}x$ on nodal vectors, where $\mathbf{L}$ is the *stiffness matrix*.

In addition taking into account the (lumped) *mass matrix* $\mathbf{M}$ with $\mathbf{M}\bar{U} \cdot \bar{V} := \int_\Omega \mathcal{I}_h(U \cdot V)\,\mathrm{d}x$ ($\mathcal{I}_h$ denotes the piecewise affine Lagrangian interpolation on $\mathcal{T}_h$), one defines in analogy to the above relation for the continuous Laplacian (1) the discrete negative Laplacian on nodal vectors $-\triangle_h \bar{U} = \mathbf{M}^{-1}\mathbf{L}\bar{U}$. Finally, based on (2) a discrete bi-Laplacian on nodal vectors is given by $\triangle_h^2 \bar{U} = \triangle_h \triangle_h \bar{U}$, which corresponds to a discrete thin plate energy $\mathcal{W}^h_{\mathrm{plate}}[U] = \int_\Omega (\mathcal{I}_h(\triangle_h(\bar{U})))^2\,\mathrm{d}x$.

In what follows, we will now translate this calculus into the geometric context.

Smooth embedded surfaces. Let $\mathcal{S} \subset \mathbb{R}^3$ be a smoothly embedded oriented surface. Let n denote the unit normal. Then the *first and second fundamental forms* on $\mathcal{S}$ are quadratic forms acting on pairs (u, v) of tangent vectors of $\mathcal{S}$, defined as

$$\mathrm{I}(u, v) := \langle u, v\rangle_{\mathbb{R}^3} \quad \text{and} \quad \mathrm{II}(u, v) := \mathrm{I}(dn(u), v)\,,$$

respectively. Here $\langle u, v\rangle_{\mathbb{R}^3}$ denotes the standard Euclidean inner product and $dn(u)$ is the directional derivative of n in the direction of u. Notice that $dn(u)$ is automatically tangential since n has constant (unit) length. The first fundamental form measures the metric of $\mathcal{S}$, while the second fundamental form, by accounting for the change of normals, measures curvature. Both first and second fundamental form are *symmetric*. The *shape operator* is the linear mapping corresponding to the second fundamental form, i.e., $\mathrm{II}(u, v) = \mathrm{I}(\mathbf{S}(u), v)$. The eigenvalues κ_1 and κ_2 are the *principal curvatures*. The *Gauß* and *mean* curvature are defined as the determinant and trace of $\mathbf{S}$, respectively.

The analogue of the the Laplacian on $\mathcal{S}$ is the Laplace–Beltrami operator $\triangle_\mathcal{S} = -\mathrm{div}_\mathcal{S}\nabla_\mathcal{S}$. Its definition requires the notions of the *gradient* of functions $f : \mathcal{S} \to \mathbb{R}$ and the *divergence* of tangential vector fields u on $\mathcal{S}$. The former is defined via $\mathrm{I}(\nabla_\mathcal{S} f, u) = df(u)$ and the latter is defined as the (negative formal) adjoint of $\nabla_\mathcal{S}$, i.e, $\int_\mathcal{S} \mathrm{I}(\nabla_\mathcal{S} f, u)da = -\int_\mathcal{S} f\mathrm{div}_\mathcal{S}(u)da$, where da denoted the area element of $\mathcal{S}$. With these definitions, one can define the Dirichlet energy of functions $\mathcal{W}_{\mathrm{Dirichlet}}[u] = \frac{1}{2}\int_\mathcal{S} |\nabla_\mathcal{S} u|^2\,\mathrm{d}a$, in perfect analogy to the flat case. Furthermore, as in the flat case, the weak formulation of the Laplace–Beltrami operator *only* requires the notion of gradient, i.e.,

$$\int_\mathcal{S} -\triangle_\mathcal{S} u\phi\,\mathrm{d}a = \partial_u \mathcal{W}_{\mathrm{Dirichlet}}[u](\phi) = \int_\mathcal{S} \nabla_\mathcal{S} u \cdot \nabla_\mathcal{S}\phi\,\mathrm{d}a \tag{3}$$

for compactly supported, smooth functions $\phi : \mathcal{S} \to \mathbb{R}$. Instead of scalar functions, we can consider vector values functions and apply differential operators to all components of the function. With a slight misuse of notation, we write x also for the vector values mapping on $\mathcal{S}$, which maps every point x onto itself. The fundamental geometric insight is that $\triangle_\mathcal{S} x = -hn$. The mean curvature vector hn is also the gradient of the area functional $\mathcal{A}[x] = \int_\mathcal{S} \mathrm{d}a$ with respect to the L^2 scalar product. Hence, the associated gradient flow of the area functional is the geometric heat equation $\partial_t x - \triangle_\mathcal{S} x = 0$. A the nonlinear counterpart of the thin plate energy on flat domains is the Willmore energy $\mathcal{W}_{\mathrm{Willmore}}[x] = \frac{1}{2}\int_\mathcal{S} h^2\,\mathrm{d}a$. The associated gradient flow is the Willmore flow $\partial_t x = (\triangle_\mathcal{S} h + h(|\mathbf{S}|^2 - \frac{1}{2}h^2))n$, where $|\mathbf{S}|$ denotes the Frobenius norm of the shape operator.

In practice, *triangle meshes* or what is known as *polyhedral surfaces* are a prevalent *discrete* representation of surfaces on a computer. In doing so, one can then, just like in the flat case, construct stiffness and mass matrices as described above. We outline this description below. One challenge is to define the second fundamental form in this case, which we also discuss below. Before doing so, we briefly recall how first and second fundamental form and the Laplace–Beltrami operator are represented in a local parameterization of $\mathcal{S}$.

Local parameterizations. Given a surface $\mathcal{S} \subset \mathbb{R}^3$, one can in general not obtain a *global parameterization*, i.e., a diffeomeorphism from a region in a plane to $\mathcal{S}$, that covers all of $\mathcal{S}$, even for simple surfaces such as the sphere. In general, one can only work with *local* parameterizations or what is known as *coordinate charts*. Working with local parameterization frequently requires to account for coordinate changes, which is a tedious task in practice. Therefore in geometry processing one prefers to work with global representations of surfaces such as triangle meshes or level set representations that do not require any particular parameterization. Moreover, the *intrinsic*, i.e., parameterization independent, formulation of differential operators is often more natural and convenient. Nonetheless, certain applications, such as texture mapping, do require local parameterizations and higher order discretizations methods such as subdivision

finite elements use local parametrization over control meshes as well. There is an extensive literature (see, e.g., [34] and references therein) on the computation of local parameterizations in particular dealing with the issue to minimize *distortion of lengths and angles*. Since distortion is unavoidable in general, this poses the question if one can at least bound the amount of introduced distortion. For the case of triangle meshes and piecewise linear continuos bijections from a polyhedral surface to a planar region, this long standing problem has recently been solved by Lipman for the case of aspect ratio distortion [58], see Figure 1. A detailed discussion of local parameterizations is beyond the scope of this paper.

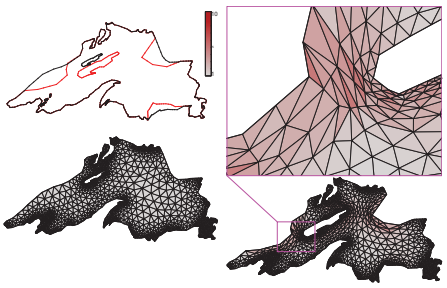

Figure 1: Bounded distortion mapping of Lake Superior. Original mesh (left) is deformed by moving island and boundaries. Bounded distortion maps (right) avoid triangle flips and high distortion. (Image courtesy of Y. Lipman.)

Returning to representations of differential operators and fundamental forms in local coordinate charts of smooth surfaces, let $x : \Omega \to \mathbb{R}^3$, $\xi \mapsto x(\xi)$ be a local parameterization of $\mathcal{S}$ defined on a parameter domain $\omega \subset \mathbb{R}^2$. Then the normal $n(x)$ is given by $n(x) = \partial_{\xi_1} x \wedge \partial_{\xi_2} x$ and the first fundamental form on the parameter domain is expressed by $g(v,w) = Dx\, v \cdot Dx\, w$, where $Dx(\xi) \in R^{3,2}$ is the Jacobian of the parameterization x. The associated matrix is $g = (g_{ij})_{i,j=1,2} = Dx^T Dx$ and its inverse $g^{-1} = (g^{ij})_{i,j=1,2}$. With the metric at hand, the integral of a function f on $x(\omega)$ is given by $\int_{x(\omega)} f \,\mathrm{d}a = \int_\omega f \circ \phi \sqrt{g}\,\mathrm{d}\xi$. For the shape operator $\mathbf{S}$ defined above we obtain from $\mathbf{S}(x) Dx\, v \cdot Dx\, w = D(n \circ x) v \cdot Dx\, w$ the representation on the parameter domain $\mathbf{S}^\omega = g^{-1} Dx^T D(n \circ x)$ as a 2×2 matrix defined on the chart. The above mentioned differential operators can as well be expressed on the chart. For the tangential gradient we obtain $\nabla_{\mathcal{S}} f(x) = Dx g^{-1} \nabla (f \circ x)$ and for the (tangential) divergence operator on $\mathcal{S}$ applied to a tangential vector field $\mathbf{v} \circ x = Dx\, v$ one deduces the parametric representation $(\mathrm{div}_{\mathcal{S}} \mathbf{v}) \circ x = \frac{1}{\sqrt{g}} \mathrm{div}_\xi (v\sqrt{g})$. Hence, the Laplace Beltrami operator is given by $\triangle_{\mathcal{S}} u = \mathrm{div}_{\mathcal{S}} \nabla_{\mathcal{S}} u = \frac{1}{\sqrt{g}} \mathrm{div}_\xi \left(g^{-1} \nabla_\xi (u \circ x) \sqrt{g} \right)$.

Triangulated surfaces. Following the discretization procedure in the flat Euclidian case, we consider a triangulated surface $\mathcal{S}_h$, e.g., the vertices of the triangulation might lie on the true surface $\mathcal{S}$ (cf. Dziuk and Elliott [30]). Then we denote by $\mathcal{V}_h$ the space of piecewise affine, continuous functions on $\mathcal{S}_h$. Again the variation of the discrete Dirichlet energy $\mathcal{W}^h_{\mathrm{Dirichlet}} = \frac{1}{2} \int_{\mathcal{S}_h} |\nabla_{\mathcal{S}_h} U|^2 \,\mathrm{d}a$ for functions $U \in \mathcal{V}_h$ leads to a quadratic form $\mathbf{L}\bar{U} \cdot \bar{V} := \int_{\mathcal{S}_h} \nabla_{\mathcal{S}_h} U \cdot \nabla_{\mathcal{S}_h} V \,\mathrm{d}a$ on nodal vectors $\bar{U}$ and $\bar{V}$ (U, V denoting the associated discrete functions in $\mathcal{V}_h$) with the stiffness matrix $\mathbf{L}$. Here, the (discrete) surface gradient $\nabla_{\mathcal{S}_h} U$ is constant on each triangle T of $\mathcal{S}_h$ and lies in the plane of the triangle, which coincides with the local tangent space. The entries of the local stiffness matrix can be computed as follows. Let us denote by v_0, v_1, v_2 the vertices of T and by e_j the *oriented* edge opposite of v_j for $j = 0, 1, 2$, then we obtain the local stiffness matrix $(\mathbf{L}_h^{\mathrm{local}})_{ij} = \frac{1}{4} e_i \cdot e_j |T|^{-1}$ and the local (lumped) mass matrix $(\mathbf{M}_h^{\mathrm{local}})_{ij} = \frac{1}{3} |T| \delta_{ij}$, which have to be assembled to there global counterparts in the usual way. Notice that $(\mathbf{L}_h^{\mathrm{local}})_{ij} = -\frac{1}{2} \cot(\alpha_{ij}))$, where $\alpha_{ij} = \angle(e_i, e_j)$ denotes the (unoriented) angle between e_i and e_j in T. The latter representation is know as the *cotangent formula* [62, 72]. Analogously, we define the mass matrix $\mathbf{M}\bar{U} \cdot \bar{V} := \int_{\mathcal{S}_h} U \cdot V \,\mathrm{d}a$ on the discrete surface $\mathcal{S}_h$. From the discrete analog of (3) we deduce the definition of the discrete Laplace Beltrami operator $-\triangle_{\mathcal{S}_h} \bar{U} = \mathbf{M}^{-1} \mathbf{L} \bar{U}$ on nodal vectors $\bar{U}$ [28]. *Convergence* of this approach for polyhedral surfaces that approximate (but do not necessarily interpolate) a smooth limit surface has been established in [45].

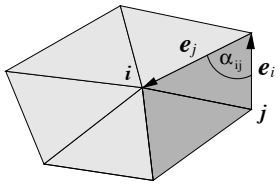

Discrete first and second fundamental forms One challenges for polyhedral surface is to provide a consistent notion of the second fundamental form, i.e., a notion that converges (in an appropriate topology or measure theoretic sense) to the second fundamental form of a smooth limit surface $\mathcal{S}$ for a sequence of polyhedral surfaces that converge to $\mathcal{S}$ in the limit of triangle refinement. In a nutshell, the difficulty is that curvatures in the smooth case involve second derivatives—a concept that does not immediately translate to

triangular meshes.

One computationally popular model for discretizing the second fundamental form is Grinspun's et al. *discrete shells* (DS) [38], which offers a discrete version of *mean curvature* associated with edges. In this view, mean curvature is measured as the dihedral angle of a an edge. It turns out that this view corresponds to a beautiful mathematical approach—the mean curvature *measure* arising form the theory of *normal cycles* [66]. In the smooth case, the theory of normal cycles is intimately connected to the method of moving frames, where one considers adapted orthonormal frames in that sense that one of the three frame vectors agrees with the unit surface normal. Using normal cycles, one not only arrives at a natural version of mean curvature for polyhedral surfaces, but also a notion of discrete Gauß curvature, associated with mesh vertices, defined as the usual *angle deficit*. This notion gives rise to a discrete version of the famous Gauß–Bonnet formula. Indeed, the power of the theory of normal cycles is that it provides a *unified* theory for the case of smooth surfaces, convex (but not necessarily smooth) bodies, and triangular meshes.

Cohen–Steiner and Morvan have used the theory of normal cycles to provide a notion of discrete shape operators for polyhedral surfaces [22, 23]. Roughly speaking, the attendant notion of second fundamental forms arises by observing that (i) there is no normal curvature along an edge and (ii) (a function of) normal curvature perpendicular to an edge is equal to the edge's dihedral angle. This notion turns out to converge in the sense of measures. A similar notion that also converges in these sense of measures, leading to a consistent notion of the Willmore functional on polyhedral surfaces, has been proposed by Hildebrandt and Polthier [43, 44]. Convergence in the sense of measures can be established since the biasing local effect of edge directions tends to decrease when averaged over large enough regions. Consequently, convergence can be shown by letting the averaging region shrink at a *much* lower rate than the refinement of triangles increases within the region. The averaging effect is much more pronounced for unstructured meshes (e.g., Dealaunay triangulation) than it is for structured meshes, see [37].

To summarize, the main idea is to *not* a priori fix a normal per edge but to initially allow for all normals that are perpendicular to their respective edge, and letting (global) bending energy minimization decide for which normal to pick. While this approach works favorably for both unstructured and structured meshes, it is still an challenging open problem to prove that this formulation is *convergent*, i.e., to show that discrete minimizers converge to smooth ones under appropriate boundary conditions.

A slight modification of these approaches leads to the following formulation of discrete second fundamental forms that mimics the construction of the smooth case, where the second fundamental form encodes the change of normals. We elaborate on this construction since it leads (i) to a formulation of second fundamental forms that is *constant* per triangle and thus (ii) to a discrete formulation of elastic energies that structurally resembles the smooth setting. For an edge $e = T \cap \tilde{T}$ between two triangles, define n_e as the normalized sum of the unit normals belonging to the triangles T and $\tilde{T}$. For boundary edges, consider the respective triangle normal, and associate n_e with the *midpoint* of e. With normals associated to edge midpoints, the (discrete) 1-form $\mathrm{d}n$ acts on line segments connecting edge midpoints. Fixing a triangle T with edges e_1, e_2, e_3 and corresponding edge normals n_1, n_2, n_3 one finds that $\mathrm{d}n(m_{ij}) = \int_{m_{ij}} \mathrm{d}n = n_j - n_i$, where m_{ij} is the line segment connecting the midpoint of e_i with that of e_j. Using the vector identity $e_k = -2m_{ij}$, where k is the complementary index to i and j in T, one accordingly defines $\mathrm{II}_{T,k} := \mathrm{II}_T(e_k, e_k) := 2(n_i - n_j) \cdot e_k$ as the action (associated with a triangle T) of the second fundamental form on the edge vector e_k. Assembling contributions over all three triangle edges and using the fact that in dimension two a symmetric quadratic form is *uniquely determined* by its action on three linearly independent vectors, leads to a discrete second fundamental form that is constant per triangle:

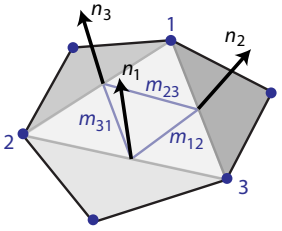

$$\mathrm{II}_T = \frac{1}{8|T|^2} \sum_{i=1}^{3} (\mathrm{II}_{T,j} + \mathrm{II}_{T,k} - \mathrm{II}_{T,i})\, t_i \otimes t_i \,, \tag{4}$$

where the indices $j = i + 1 \pmod 3$ and $k = i + 2 \pmod 3$ refer to the cyclic ordering of edges of T, $\otimes$ denotes the outer product, t_i is the result of clockwise rotating edge e_i by $\pi/2$ in the plane of T, and $|T|$ denotes the area of T.

Repeating the above construction for the first fundamental form, one defines $\mathrm{I}_{T,k} := \mathrm{I}_T(e_k, e_k) := |e_k|^2$

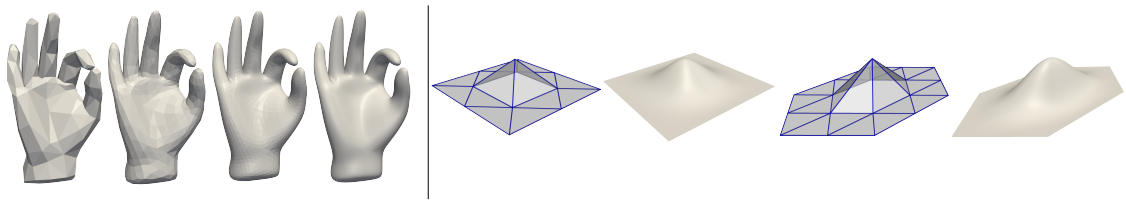

Figure 2: The first two subdivision steps starting from the control mesh and the resulting limit surface are visualized (left). Furthermore, subdivision bases functions are rendered for a node with valence 4 (left) and 6 (right), respectively (Image courtesy of R. Perl).

and thus

$$\mathrm{I}_T = \frac{1}{8|T|^2} \sum_{i=1}^{3} (\mathrm{I}_{T,j} + \mathrm{I}_{T,k} - \mathrm{I}_{T,i})\, t_i \otimes t_i \,. \tag{5}$$

In the computational mechanics community this formulation is known as the *constant strain triangle*.

Subdivision Finite Elements. As long as smooth surfaces are approximated by a triangular mesh one is restricted to the discretization of PDEs and variational problems on the surface via piecewise affine continuous finite elements as exposed above. Alternatively, a triangular mesh could be considered as a control mesh of a much smoother surface approximation. Such approximation can be generation via an iterative refinement of the control mesh using a subdivision scheme. Starting from the control mesh a single subdivision step consists of a mesh refinement and then every vertex of the refined mesh is replaced by a weighted sum of itself and all adjacent vertices with a usually fixed set of weights. In their pioneering paper Catmull and Clark [17] proposed such a scheme for quadrilateral meshes. A subdivision method for triangular grids based on a successive refinement of every triangle into four congruent triangles was presented by Loop [61] (cf. Fig. 2). If restricted to control patches with valence 6, i.e., every vertex of a triangle is shared by 6 adjacent triangles forming the patch associated with the triangle, Loop's subdivision scheme is equivalent to the construction of bi-variate box-splines and thus the limit surface is C^2 and in particular it can be constructed explicitly from the vertex positions of a local patch of the control mesh. For general meshes with different valences the limit surface is still $C^1 \cup H^2$ and thus for instance the mean curvature is still in L^2. Mandal, Qin, and Vemuri [63] used subdivision surfaces for the dynamic surface modeling via spring forces attached to the control points. Subdivision schemes are not only useful for the smooth approximation of surfaces using a local parametrization over flattened control meshes. They can be applied to assigned general scalar values or vectors to the control vertices to describe smooth scalar or vector valued functions. In particular the conforming finite element approximation of geometric functionals involving second derivatives is possible, e.g., the discrete weak definition of the geometric Laplacian $\triangle^2_{\mathcal{S}}$ on the limit surface $\mathcal{S}_h^\infty$ of a control mesh $\mathcal{S}_h$ is given by

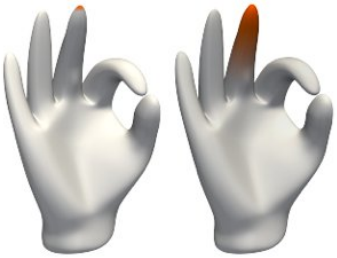

Figure 3: The equation $\triangle^2_{\mathcal{S}} u + u = f$ is solved on a surface using subdivision finite elements based on Loop subdivision with color coded f (left) and u (right) (Image courtesy of R. Perl).

$$\int_{\mathcal{S}_h^\infty} \triangle^2_{\mathcal{S}_h^\infty} U \; \Phi \, \mathrm{d}a = \int_{\mathcal{S}_h^\infty} \triangle_{\mathcal{S}_h^\infty} U \, \triangle_{\mathcal{S}_h^\infty} \Phi \, \mathrm{d}a \tag{6}$$

for a subdivision finite element function U and all subdivision finite element test functions Φ (cf. (2)). Figure 3 shows the result for solving the equation $\triangle^2_{\mathcal{S}} u + u = f$ on a surface using subdivision finite elements. Cirak, Ortiz, and Schröder [21] investigated a thin-shell analysis using subdivision finite elements to describe the surface geometry and to compute smooth displacement field in a conforming finite element

Galerkin approach. They applied this approach to the Kirchhoff-Love theory of thin shells and used Loop's subdivision scheme. In [39] Grinspun, Krysl, and Schröder have used a base function oriented adaptive strategy in order to implement an adaptive subdivision finite element scheme. For the assembly of finite element matrices on valence 6 control patches the explicit spline representation can be retrieved and the computations are performed on the control mesh as a parameter domain of the actual subdivision limit surface. On triangles of non valence 6 patches, one uses an appropriate quadrature formula with only interior quadrature points and applies the subdivision scheme recursively until all quadrature points lye in a valence 6 patch and thus the explicit representation again applies (cf. Fig. 2).

Implicit surfaces / level sets. Now we consider surfaces $\mathcal{S}$ described as level sets of a function w, i.e., $\mathcal{S}^c = [w = c] := \{x \in \mathbb{R}^3 \,|\, w(x) = c\}$, where we assume $w : \mathbb{R}^3 \to \mathbb{R}$ to be smooth with $\nabla w \neq 0$ on $\mathcal{S}^c$. Then, the normal is given as $n = |\nabla w|^{-1}\nabla w$ and for the shape operator (trivially extended to a 3×3 matrix) we obtain $\mathbf{S}^{\text{ext}} = DnP = D(|\nabla w|^{-1}\nabla w)P = |\nabla w|^{-1}PD^2wP$, where $P(x) = P[w](x) := \mathbb{1} - n(x)n(x)^T$ denotes the projection onto the tangent space. Thus for the mean curvature at the point x on $\mathcal{S}_{w(x)}$ one gets $h = \operatorname{tr}(DnP) = \operatorname{div} n$. The tangential gradient of a function $f : \mathbb{R}^3 \to \mathbb{R}$ and the tangential divergence of a vector field $v : \mathbb{R}^3 \to \mathbb{R}^3$ are given by $\nabla_{\mathcal{S}^c} f = P\nabla f$ and $\operatorname{div}_{\mathcal{S}^c} v = (P\nabla)\cdot v$, respectively. In the implicit surface case, there is generically no integration over a single surface, instead we can integrate over a bundle of surfaces $[c^- < w < c^+] = \bigcup_{c\in\{c^-,c^+\}} \mathcal{S}^c$ and the coarea formula implies

$$\int_{c^-}^{c^+} \int_{\mathcal{S}^c} f \,\mathrm{d}a \,\mathrm{d}c = \int_{[c^-<w<c^+]} f|\nabla w| \,\mathrm{d}x$$

with $|\nabla w|$ representing the density of level sets. Thus, the Dirichlet energy of a function $u : \mathbb{R}^3 \to \mathbb{R}$ integrated over a bundle of surfaces $\{\mathcal{S}_c \,|\, c \in \{c^-, c^+\}\}$ can be computed by

$$\mathcal{W}^{c^-,c^+}_{\text{Dirichlet}}[u] = \frac{1}{2}\int_{[c^-<w<c^+]} (P\nabla u)\cdot\nabla u \,|\nabla w| \,\mathrm{d}x\,.$$

Likewise, one obtains the weak definition for the Laplace Beltrami operator

$$\int_{[c^-<w<c^+]} -\triangle_{\mathcal{S}_w} u\,\phi\,|\nabla w| \,\mathrm{d}x = \partial_u \mathcal{W}^{c^-,c^+}_{\text{Dirichlet}}[u](\phi) = \int_{[c^-<w<c^+]} (P\nabla u)\cdot\nabla\phi\,|\nabla w| \,\mathrm{d}x \tag{7}$$

for compactly supported, smooth functions $\phi : \mathbb{R}^3 \to \mathbb{R}$ from which the explicit representation $\triangle_{\mathcal{S}_w} u = \frac{1}{|\nabla w|}\operatorname{div}(|\nabla w| P\nabla u)$ is deduced for smooth u. Taking into account the level set equation

$$\partial_t w + |\nabla w| v = 0\,,$$

which describes the evolution of the implicit surfaces $\mathcal{S}_w$ with a speed v in direction of the normal field $n = |\nabla w|^{-1}\nabla w$, we obtain for the surface evolution of mean curvature motion $\partial_t w + |\nabla w|\operatorname{div}(|\nabla w|^{-1}\nabla w) = 0$. Finally, the Willmore energy integrated over a bundle of surfaces $\{\mathcal{S}_c \,|\, c \in \{c^-, c^+\}\}$ is given by

$$\mathcal{W}^{c^-,c^+}_{\text{Willmore}}[u] = \frac{1}{2}\int_{[c^-<w<c^+]} \operatorname{div}(|\nabla w|^{-1}\nabla w)^2 |\nabla w| \,\mathrm{d}x\,.$$

Let us emphasize that in the level set formulation the PDE problems on different surfaces $\mathcal{S}_c$ and $\mathcal{S}_{\tilde{c}}$ are still decoupled even though one integrates over bundles of surfaces.

Level set Finite Elements. Discretizing level set equations is based on a triangulation $\mathcal{T}_h$ of the computational domain $D \subset \mathbb{R}^3$, where h again denotes the grid size. Now, let $\mathcal{V}_h$ denote the space of piecewise affine, continuous functions on this 3D triangulation. Then, given a discrete level set function $W \in \mathcal{V}_h$ the stiffness and matrix matrix are defined by $\mathbf{L}\bar{U}\cdot\bar{V} = \int_D P[W]\nabla U\cdot\nabla V|\nabla W| \,\mathrm{d}x$ and $\mathbf{M}\bar{U}\cdot\bar{V} = \int_D \mathcal{I}_h(UV)|\nabla W| \,\mathrm{d}x$, respectively. Here, $\mathcal{I}_h$ is the affine Lagrangian interpolation on $\mathcal{T}_h$ and $\bar{U}$, $\bar{V}$ are again the nodal vectors corresponding to discrete functions U, $V \in \mathcal{V}_h$. Hence, one obtains for the discrete Laplacian on a bundle of surfaces in the level set context $-\triangle_h\bar{U} = \mathbf{M}^{-1}\mathbf{L}\bar{U}$. The finite element analysis of PDEs on level sets is discussed in [31]. For the general concepts of the level set method we

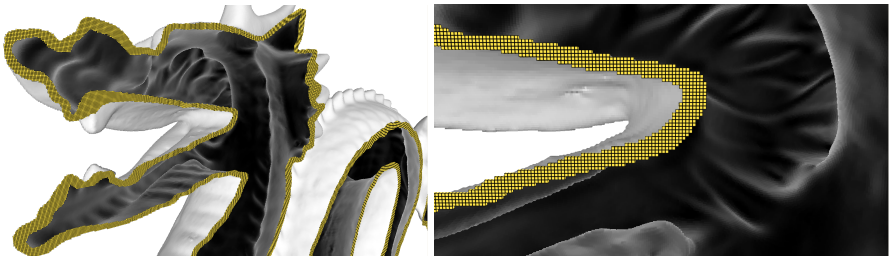

Figure 4: The level set representation of a complex dragon model on a thin narrow band is shown using a cutting plane intersecting the dragon surface with underlying resolution $1986 \times 1323 \times 1104$). On the right a zoom is displayed and on the plane the slices of all cells intersecting the narrow band are drawn.

Figure 5: *Poisson Surface Reconstruction* proposed in [54] for high quality surface reconstruction from oriented points clouds. (Image courtesy of M. Kazhdan.)

refer to the textbooks by Sethian [79] or Osher and Fedkiw [70]. If one is primarily interested in a PDE on a single surface or its geometric evolution, a narrow band approach is advisable, where one chooses as a computational domain only a small neighbourhood of the surface of interest. The efficient encoding of narrow bands on very high-resolution grids can be done with suitable hierarchical sparse grid structures as proposed in [47]. Fig. 4 depicts the narrow band geometry for a large scale surface model.

Point clouds. In practice, e.g., when using a 3D scanner, surfaces are sometimes given as raw point data—without any connectivity or mesh structure. There exist numerous algorithms for reconstructing a surface from a point set, such as discrete computational geometry methods that provide triangulated surfaces, e.g., by using Delaunay triangulations, Voronoi diagrams, or alpha complexes, [2, 12, 32]. Other approaches base reconstruction on implicit surfaces, e.g., by using radial basis functions (RBFs) or moving least squares [68,88]. If the points in the point set are additionally *oriented*, i.e., come with a unit normal per point, then one can use a PDE-based approach for reconstruction by solving a simple Poisson problem [54] (see Figure 5). Laplacians on point clouds have been constructed using short-time properties of the heat kernel, see [7].

Smooth and Discrete Exterior Calculus. On smooth or polyhedral manifolds there usually does not exist a global parameterization. Therefore, it is desirable to express differential operators in an *invariant* manner, i.e., independent of any particular local parameterization. In the smooth setting, one prevalent choice is the language of *exterior calculus* on abstract manifolds $\mathcal{M}$. We will later apply this to two dimensional surfaces $\mathcal{S}$. Exterior calculus starts with the notion of smooth differentiable k-forms. Restricted to a point $x \in \mathcal{M}$, a k-form ω_x is an alternating multilinear for, i.e., $\omega_x \in \mathrm{Alt}^k(T_x\mathcal{M})$, where $T_x\mathcal{M}$ is the tangent space at x. Denoting the space of smooth k-forms by $\Omega^k(\mathcal{M})$, one then considers the *exterior derivative* $d : \Omega^k(\mathcal{M}) \to \Omega^{k+1}(\mathcal{M})$ and the *wedge product* $\wedge : \Omega^k(\mathcal{M}) \times \Omega^l(\mathcal{M}) \to \Omega^{k+l}(\mathcal{M})$. Since this is not the place to elaborate on the various intruiging properties of these operators, we content ourselves with pointing out that exterior differentiation is a generalisation of vector calculus in $\mathbb{R}^3$, i.e., of operators such as grad and div; notice, however, that exterior differentiation *does not* require the notion of a metric.

Likewise, the wedge product is a generalization of the cross product on $\mathbb{R}^3$. If $\mathcal{M}$ is additionally equipped with a Riemannian metric g, then this metric naturally gives rise to norms of k-forms and to measuring volume on $\mathcal{M}$. Indeed, equipped with a Riemannian metric, one can define the *volume form* $dvol_g$ and the *Hodge star operator* $\star : \Omega^k(\mathcal{M}) \to \Omega^{n-k}(\mathcal{M})$, where $n = \dim(\mathcal{M})$. These objects are related through $\omega \wedge (\star\omega) = \|\omega\|^2 dvol_g$. On two-dimensional Riemannian manifolds—i.e., the case of interest here—the Hodge star on 1-forms simply corresponds to a rotation by $\pi/2$, which, if we identify vectors and 1-forms, reveals a simple geometric meaning of this operator. Likewise, the Hodge star of a 0-form (i.e., a smooth function) f yields $\star f = f\, dvol_g$, i.e., a f-weighted volume form. A Riemannian metric additionally gives rise to an L^2 inner product of k-forms, which, using the Hode star, can conveniently be expressed as $(\omega, \eta)_{L^2} = \int_{\mathcal{M}} \omega \wedge (\star\eta)$ (whenever this integral is defined). The L^2 inner product then gives rise to the *codifferential* $\delta : \Omega^k(\mathcal{M}) \to \Omega^{k-1}(\mathcal{M})$ as the (formal) adjoint of d, i.e, $(d\omega, \eta)_{L^2} = (\omega, \delta\eta)_{L^2}$. In terms of the Hodge star, the codifferential can be expressed as $\delta = (-1)^{n(k+1)+1} \star d\star$. Finally, the Hodge–Laplace operator on k-forms is defined as $\triangle_{\mathcal{M}} = \delta d + d\delta$. Notice that, just as in the Euclidean case, $\triangle_{\mathcal{M}}$ is (formally) self-adjoint with respect to $(\cdot,\cdot)_{L^2}$.

Mimicking the above on simplicial manifolds, Desbrun, Hirani, Leok, and Marsden introduced *discrete exterior calculus* (DEC) [26]. In this case, k-forms are naturally associated with simplicial k-cochains C^k, i.e., duals to the space of simplicial k-chains C_k (which in turn are defined as formal $\mathbb{R}$–linear combinations of k-simplices). The role of exterior differentiation is the canonically taken by the simplicial coboundary operator $\delta : C^k \to C^{k+1}$. For constructing a discrete Hode star operator, Hirani et al. consider the *circumcentric* dual complex K^* of a simplicial complex K. This dual construction leads them to a discrete Hodge star operator taking primal cochains in K to dual cochains in K^* and vice-versa. properly weighted) dual of $* : C_k \to C_{n-k}$. In analogy to the smooth setting, one then defines the dual of the coboundary operator by $\delta^* := (-1)^{n(k+1)+1} \star d\star$, mapping C^k to C^{k-1}, and the discrete Laplacian as $\triangle_h := \delta^*\delta + \delta\delta^*$. For the Laplacian acting on 0-cochains (i.e., discrete functions defined at vertices), one recovers the cotangent operator from the FE setting described above—with the slight modification that the attendant mass matrix in the DEC construction is diagonal, corresponding to a *lumped* mass matrix in the FE world. An alternate approach for defining δ^* (and therefore $\triangle_h$) is, as the the smooth case, to consider L^2-inner products $(\cdot,\cdot)_{L^2}$ on k-cochains and defining δ^* as the dual of δ with respect to this inner product. Indeed, an appropriate choice of inner products leads to the same δ^* (and hence to the same Laplacian) as the one arising from the discrete Hodge star. Using inner products additionally leads to a construction of Laplacians on surfaces with general polygonal (not necessarily triangular) faces [1].

In a similar spirit, Arnold, Falk, and Winther have developed *Finite Element Exterior Calculus* (FEEC) [3]. One of the main approaches there is to avoid an explicit construction of the formal dual of d and instead use a weak formulation on suitable finite element spaces, i.e., instead of solving $\delta\sigma = \omega$, one considers the weak formulation $(\omega, \tau) = (\sigma, d\tau)$, leading to a *mixed* formulation for the Hodge Laplacian. For piecewise linear functions, DEC and FEEC share various similarities—the main difference being that DEC works with a diagonal (lumped) mass matrix, whereas FEEC does not.

The approaches of DEC and FEEC have in common a certain trend to consider *structure preserving* discretizations. For the case of partial differential equations, the benefit of structure preservation appears to be the resulting *stable* discretizations. A similar observation applies to geometry: For example, the structure of special surfaces, such as minimal surfaces, surfaces of constant mean curvature or surfaces of constant negative Gauß curvature, is governed by nonlinear PDEs. The field of *Discrete Differential Geometry* (DDG) searches for discretizations that preserve the underlying structure of the smooth case. For readers interested in theory and applications of DDG to geometry processing, we refer to [9–11].

Finally, with regards to structure preservation, there are certain limits to what one is able to achieve in the discrete case when trying to mimic *all* properties of the smooth setting. E.g., for the Laplace–Beltrami operator it has been shown [91] that there exists no discrete Laplacian that possesses *all* properties of the smooth one. Indeed, the main culprit in the discrete setup is to construct Laplacians that are convergent *and* satisfy the *maximum principle*.

3 Surface Fairing and Surface Flows

In physics, diffusion is known as a process that equilibrates spatial variations in concentration. For a (noisy) concentration u_0 on a domain $\Omega \subset \mathbb{R}^2$ the heat equation $\partial_t u - \triangle u = 0$ with natural boundary conditions describes a scale of representations $\{u(t)\}_{t\in\mathbb{R}^+}$ of the initial values u, which gets successively

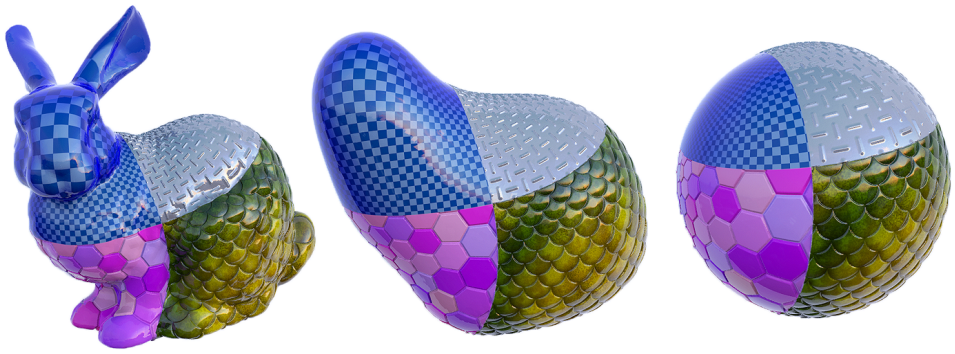

Figure 7: Willmore flow on triangulated meshes when working in *curvature space*, see [24]. (Image courtesy of K. Crane.)

coarser for $t \to \infty$. It is well-known that for $\Omega = \mathbb{R}^2$ the solution u coincides with a Gaussian filtering of the initial data, i.e., $u(\sigma^2/2) = G_\sigma * u_0$ with $G_\sigma(x) = (2\pi\sigma^2)^{-1} e^{-x^2/(2\sigma^2)}$. Here σ denotes the filter width or standard deviation. It is near at hand to ask for analogous strategies for the fairing of perturbed and noisy surface geometries $\mathcal{S}_0$ and is naturally lead to mean curvature motion $\partial_t x - \triangle_{\mathcal{S}} x = 0$ using the fact that $hn = -\triangle_{\mathcal{S}} x$. Discretizing this evolution problem explicitly in space one obtains a first simple fairing scheme $x^{n+1} = x^n - \tau \triangle_{\mathcal{S}^n} x^n$ with x^n denoting the identity on the surface $\mathcal{S}^n$. To speed up this scheme and improve its robustness Desbrun et al. [27] considered the implicit time discretization $x^{n+1} = x^n - \tau \triangle_{\mathcal{S}^n} x^{n+1}$ similar to the one proposed by Dziuk [29], where the Laplace-Beltrami operator is still evaluated at the old time step [27] (cf. Fig. 6). Mean curvature motion model is known as the L^2 gradient flow of the surface area $\mathcal{A}(\mathcal{S})$ and $\frac{d}{dt}\mathcal{A}(\mathcal{S}(t)) = -\int_{\mathcal{S}(t)} h^2 \, da$ (cf. [49]), which is one indication for the strong regularizing effect of mean curvature motion. Unfortunately mean curvature motion not only decreases the geometric noise due to imprecise aquisition but it also smoothes out geometric features such as edges and corners. Hence, a models is required which improve a simple high pass filtering. Here image processing methodology can be used. Perona and Malik [71] proposed a nonlinear diffusion method based on the PDE $\partial_t u - \operatorname{div}(a(|\nabla u|)\nabla u) = 0$ with a diffusion coefficient $a(s) = \left(1 + \frac{s^2}{\lambda^2}\right)^{-1}$ which suppresses diffusion in areas of high gradients with an edge classifier constant $\lambda > 0$, which leads to sharpening by backward diffusion whenever $|\nabla u| \geq \lambda$, whereas the image is smoothed elsewhere by forward diffusion (cf. [53]). To overcome the ill-posedness of the original Perona and Malik model Catté et al. [18] proposed a regularization method, where the diffusion coefficient $a(\cdot)$ is evaluated on a prefiltered image intensity $u_\sigma = G_\sigma * u$. Weickert [92] improved this method using an anisotropic diffusion tensor, where the Perona Malik type diffusion is concentrated in the gradient direction of the prefiltered image u_σ. This implies an additional tangential smoothing along edges. This approach can be transferred to geometry processing. Thereby edges and corners are classified using the shape operator $\mathbf{S}_{\mathcal{S}_\sigma}$ of a prefiltered surface $\mathcal{S}_\sigma$. Close to an edge the principal direction of curvature corresponding to the dominant principal curvature points in the direction orthogonal to the edge.

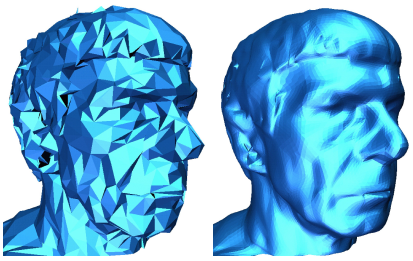

Figure 6: The smoothing of an initial surface (left) by time discrete mean curvature motion (right) using the implicit fairing approach by Desbrun et al. [27] (Image courtesy of M. Desbrun).

Triangulated surfaces. In case of an explicit surface representation one obtains the evolution problem

$$\partial_t x - \operatorname{div}_{\mathcal{S}}(A^\sigma \nabla_{\mathcal{S}} x) = 0\,,$$

where $A^\sigma = a(\mathbf{S}_\sigma)$ is a diffusion tensor, which is diagonal with respect to the orthonormal basis of principal curvature directions v_σ^i on $\mathcal{S}_\sigma$, i.e., $A^\sigma v_\sigma^i \cdot v_\sigma^j = \delta_{ij} a(\kappa_\sigma^i)$ with $\kappa^{i,\sigma}$ denoting the

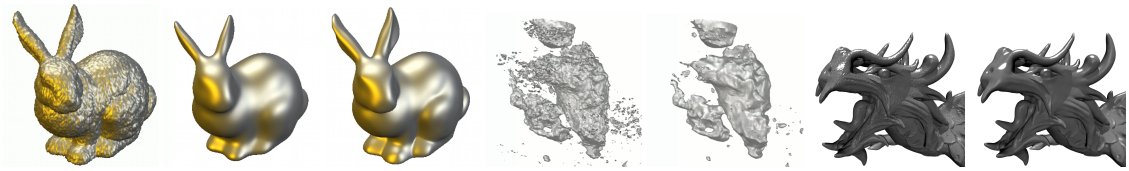

Figure 9: Mean curvature motion fairing is compared with anisotropic geometric diffusion for a noisy triangular surface (left), to a level set of a 3D echocardiographical image of a ventricle of the human heart (middle) and to a a noisy level set model of a dragon using a narrow band approach for high resolution surfaces (right).

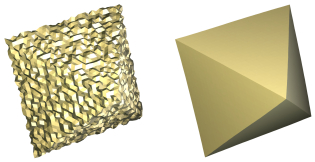

Figure 8: Result of the discrete geometry approach to anisotropic diffusion (right) for a noisy octahedron (left) (Image courtesy of K. Hildebrandt).

corresponding principal curvatures. Following the general discretization procedure outlined in Section 2 the method can be discretized on triangular meshes using affine finite elements. An alternative approach has been investigated by Hildebrandt and Polthier [42]. They used a discrete shape operator on triangular meshes to model a discrete anisotropic diffusion process directly on the triangular mesh as demonstrated in Fig. 8. A subdivision finite element implementation of anisotropic diffusion was proposed by Bajaj and Xu [4]. Fig. 9 shows results of the anisotropic diffusion method. By working in *curvature space* instead of in position space, Crane et al. [24] have recently proposed a formulation of Willmore flow for triangulated surfaces that allows for very large time steps while preserving the quality of the input mesh, see Figure 7. Their approach can also be used for highly efficient surface fairing while preserving mesh quality.

Implicit surfaces / level sets. This approach can be adapted to the processing of level set surfaces. Let us denote by $w_0 : D \to \mathbb{R}$ the implicit representation of a surface $\mathcal{S}_0$ ($\mathcal{S}_0 = [w_0 = 0]$). Then one asks for a family $\{w(t)\}_{t>0}$ of denoised level set functions which solves the anisotropic diffusion problem

$$\partial_t \phi - |\nabla \phi| \operatorname{div}\left(A^\sigma \frac{\nabla \phi}{|\nabla \phi|} \right) = 0$$

with natural (no flux) boundary conditions on ∂D and initial data $w(0) = w_0$. In this case $A^\sigma = a(\mathbf{S}_\sigma^{\text{ext}}) - n_\sigma \otimes n_\sigma$ with $\mathbf{S}_\sigma^{\text{ext}}$ being the extended shape operator of a smoothed representation of the level set function w and n_σ the associated smoothed normal. In particular at $x \in D$, one might consider the L^2 projection of w on a ball $B_\sigma(x)$ onto the space of quadratic polynomials on $\mathbb{R}^3$. Finally, for the spatial discretization can be achieved via finite element discretization of the computational domain (cf. Section 2). Figure 9 shows the application of the anisotropic diffusion method in the context of 3D medical imaging [73]. Furthermore the method is applied to a very detailed surface, where the implementation is based on hierarchical sparse narrow band structures [69].

4 Elastic Energy of Thin Shells

When surfaces are deformed, different types of deformations are observed. To analyze those from a mechanical perspective we view surfaces as thin elastic shells $\mathcal{S}^\delta$, defined as δ thick layer of material around the center surface $\mathcal{S}$, i.e.,

$$\mathcal{S}^\delta = \{x + sn(x) \,|\, x \in \mathcal{S},\ -\delta < s < \delta\}\,. \tag{8}$$

For small thickness δ one mainly distinguishes (cf. Fig. 10):

(i) tangential distortion caused by in-layer (tangential) shear, compression, or expansion,

(ii) (transversal) shear caused by normal bending.

With respect to applications in geometry processing one is not interested in the volumetric deformation of the thin layer but in the effective behavior of the thin shell in the asymptotical limit for vanishing thickness. In this limit the two types of deformation energies scale differently with respect to the thickness δ and the associated deformation energies solely refer to the geometry of the undeformed and deformed surface $\mathcal{S}$.

Tangential distortion energy. Let $\mathcal{S}$ be locally parametrized over a parameter domain ω with a parametrization x (cf. Section 2) and let $g^{\mathcal{S}}$ denote the associated metric on ω. If $\phi : \mathcal{S} \to \mathbb{R}^3$ denotes a smooth deformation of $\mathcal{S}$, $\psi = \phi \circ x : \omega \to \mathbb{R}^3$ is a local parametrization of the deformed surface $\phi(\mathcal{S})$. The metric of $\mathcal{S}$ and $\phi(\mathcal{S})$ are given by $g^{\mathcal{S}} = Dx^T Dx$ and $g^{\phi(\mathcal{S})} = D\psi^T D\psi$, respectively. For a deformation $\phi : \mathbb{R}^3 \to \mathbb{R}^3$ the mechanically relevant deformation tensor is the frame indifferent Cauchy Green strain tensor $\mathbf{G}[\phi] = D\phi^T D\phi$, which represents the pull back of the Euclidian product from the deformed configuration in the undeformed (reference) configuration. For a deformation ϕ from a surface $\mathcal{S}$ to a deformed surface $\phi(\mathcal{S})$ the parametric (tangential) Cauchy Green strain tensor describing the tangential distortion is given by the pull back $\mathbf{G}^{\omega}[\phi]$ of the metric $g^{\phi(\mathcal{S})}$ of the deformed configuration in the metric $g^{\mathcal{S}}$ of the undeformed configuration, i.e., $g^{\mathcal{S}}(\mathbf{G}^{\omega}[\phi]v, w) = g^{\phi(\mathcal{S})}(v, w)$ for all v, $w \in \mathbb{R}^2$ and hence $\mathbf{G}^{\omega}[\phi] = (g^{\mathcal{S}})^{-1} g^{\phi(\mathcal{S})}$. For the special case, were two surfaces $\mathcal{S}_A$ and $\mathcal{S}_B$ are parametrized over two parameter domains ω_A and ω_B via parameterizations x_A and x_B, respectively, a deformation $\phi : \mathcal{S}_A \to \mathcal{S}_B$ might be described by a deformation $\phi^{\omega} : \omega_A \to \omega_B$ with $\phi = x_B \circ \phi^{\omega} \circ x_A^{-1}$. In this case, by definition $g^A(\mathbf{G}^{\omega}[\phi]v, w) = (g^B \circ \phi^{\omega})(D\phi^{\omega} v, D\phi^{\omega} w)$ and thus $\mathbf{G}^{\omega}[\phi] = (g^A)^{-1}(D\phi^{\omega})^T (g^B \circ \phi^{\omega}) D\phi^{\omega}$.

One can also express the (tangential) Cauchy Green strain tensor directly on the embedded tangent space and obtains $\mathbf{G}[\phi] = \mathbf{G}^{\omega}[\phi^{\omega}] = (D_{\mathcal{S}}\phi)^T D_{\mathcal{S}}\phi$, where $D_{\mathcal{S}}\phi$ is the linear mapping from the tangent space to $\mathbb{R}^3$ with $D_{\mathcal{S}}\phi(v) = \frac{\mathrm{d}}{\mathrm{d}t}\phi(c(t))\big|_{t=0}$ for any smooth curve c on $\mathcal{S}$ with $c(0) = x$ and $\dot{c}(0) = v$. Here, $D_{\mathcal{S}}$ coincides with total variation operator d from the Discrete External Calculus. In what follows we will use both notations, where $D_{\mathcal{S}}$ is the widespread notation in mechanics and d is the usual notation in geometry. Alternatively, $D_{\mathcal{S}}\phi = D\phi^{\mathrm{ext}} P$ holds for an extension ϕ^{ext} of the deformation ϕ on the neighborhood of the surface $\mathcal{S}$ and the tangential projection P. In the coordinate system $\{\partial_{\xi_1} x,\, \partial_{\xi_2} x,\, n\}$ the two different Cauchy Green strain tensors are related by

$$\mathbf{G}[\phi] = P(D\phi^{\mathrm{ext}})^T D\phi^{\mathrm{ext}} P = \begin{pmatrix} \mathbf{G}^{\omega}[\phi] & 0 \\ 0 & 0 \end{pmatrix}.$$

The impact of tangential distortion on the elastic energy is reflected by a tangential distortion (membrane) energy, the energy density of which depends solely on the Cauchy Green strain tensor and it scales scales linear in the thickness δ, i.e.,

$$\mathcal{W}_{\mathrm{mem}}[\phi] = \delta \int_{\mathcal{S}} W_{\mathrm{mem}}(\mathbf{G}[\phi] + n_A \otimes n_A)\, \mathrm{d}a\,. \tag{9}$$

Bending energy. Bending of thin shells is described in terms of the change is the variation of the normal on the surface. Hence, to quantify the bending one aims at comparing the shape operator on the deformed surface $\mathbf{S}_{\phi(\mathcal{S})}(x)$ with the shape operator on the undeformed surface $\mathbf{S}_{\mathcal{S}}(x)$. One obtains for a parametric surface

$$\mathbf{S}_{\mathrm{rel}}(x) Dx\, v \cdot Dx\, w = D(n^{\phi} \circ \phi \circ x)v \cdot D(\phi \circ x)w - D(n \circ x)v \cdot Dx\, w$$

and thus $Dx^T \mathbf{S}_{\mathrm{rel}}(x) Dx = D(\phi \circ x)^T D(n^{\phi} \circ \phi \circ x) - Dx^T D(n \circ x)$. On the parameter domain one obtains the representation $\mathbf{S}^{\omega}_{\mathrm{rel}}(x) = g^{-1}(D(\phi \circ x)^T D(n^{\phi} \circ \phi \circ x) - Dx^T D(n \circ x))$ for the relative shape operator. For implicit surfaces the (extended) relativ shape operator is given by

$$\mathbf{S}^{\mathrm{ext}}_{\mathrm{rel}}(x) = (D_{\mathcal{S}}\phi(x))^T \mathbf{S}^{\mathrm{ext}}_{\phi(\mathcal{S})}(\phi(x)) D_{\mathcal{S}}\phi(x) - \mathbf{S}^{\mathrm{ext}}_{\mathcal{S}}(x)\,.$$

Is ϕ is locally isometric, i.e., $\mathbf{G}[\phi] = \mathbb{1}$, then $D_{\mathcal{S}}\phi$ is orthogonal and $\mathbf{S}_{\mathrm{rel}}$ measure pure bending, otherwise, the above defined relative shape operator reflects also tangential distortion. A suitable bending energy is then given by

$$\mathcal{W}_{\mathrm{bend}}[\phi] = \delta^3 \int_{\mathcal{S}} W_{\mathrm{bend}}(\mathbf{S}_{\mathrm{rel}}(x))\, \mathrm{d}a\,, \tag{10}$$

Here the scaling factor δ^3 reflects the fact that bending is a second order term in the expansion of the volumetric elastic energy and the integration volume is δ thick. A simple model for the energy density

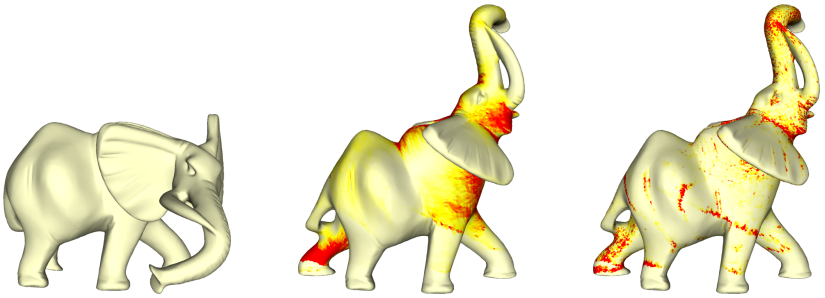

Figure 10: Two predominant types of deformation for a surface model (left) in the deformed configuration using a discrete geometry approach on triangular meshes: tangential stretching of shearing (middle, 0 0.001) and normal bending (right, 0 0.1).

is $W_{\mathrm{bend}}(\mathbf{S}_{\mathrm{rel}}(x)) = |\mathbf{S}^{\mathrm{ext}}_{\mathrm{rel}}(x)|^2$, where $|A|$ denotes the Frobenius norm of the matrix A. Notice that $W_{\mathrm{bend}}(\mathbf{S}_{\mathrm{rel}}(x))$ takes into account the full change of shape operator, not only the change of their traces (i.e., mean curvatures)—to the effect that changes of bending *directions* can be accounted for appropriately appropriately (cf. Fig. 10).

5 Deformation and Physical Simulation of Thin Shells

A designer may want to edit a shape by manually deforming only parts of a given larger rest shape, e.g., the hands, legs, or arms of a humanoid figure, through the use of handles. In this setup, the remaining parts of the shape are required to follow the deformation of the parts prescribed by the user in a plausible way. One approach for tackling this problem is through minimizing the elastic deformation energy of the deformed with respect to the undeformed shape while respecting the boundary conditions provided by the user. However, in practice, due to the nonlinearity of elastic energy, this approach is often too costly for an interactive editing session when working with detailed and complex geometries. Therefore, many deformation approaches take the route of compromise: interactive response is established at the price of sacrificing physical accuracy. Notice, though, that in graphics, the quality of a deformation method might be a evaluated by its plausibility rather than its physical accuracy. We come back to this point when discussing physical simulations, where an eyeball metric is often no longer admissible.

Isometric deformations of thin plates and quadratic bending energies For various thin elastic materials the gradient of membrane energy is usually large relative to the gradient of bending energy, since many materials tend to resist stretching more than bending [5, 87]. Therefore, one may often assume that deformations nearly preserve the metric, i.e., are *nearly isometric*. Observe that if ϕ induces an isometry and the undeformed state is *planar* in its rest state (i.e., it is a thin plate), then the first fundamental form remains unchanged and hence $\mathcal{W}_{\mathrm{mem}}[\phi] = 0$. Restricting to the case of the Frobenius norm for bending energy density, i.e., $W_{\mathrm{bend}}(A) = |A|^2$, one additionally obtains that isometric deformations of thin plates yield

$$\mathcal{W}_{\mathrm{bend}}[\phi] = \int_{\mathcal{S}} |\mathrm{Hess}[\phi]|^2 da\,,$$

where $\mathrm{Hess} = \nabla d$ denotes the Hessian of the *undeformed* surface with respect to its Riemannian metric, i.e., ∇ denotes covariant differentiation (for the ∞ dimensional analogue on shape space see Section 8). This is due to the fact that (i) the shape operator of the undeformed surface vanishes identically, thus (ii) the relative shape operator is equal to the (pullback of) the shape operator of the deformed surface, and (iii) for isometric deformations the second fundamental form of the deformed surface is related to the Hessian via $\mathrm{II}^{\phi}(v,w)n^{\phi} = \mathrm{Hess}(\phi)(v,w)$, where n^{ϕ} is the normal of the deformed surface. Notice that this implies that bending energy is *quadratic* in the displacement ϕ without any further assumptions or simplifications. This observation, when applied to the trace of the Hessian (i.e., the Laplace–Beltrami operator $\triangle$), has been

used in [90] to accelerate physical simulations. Notice that $|\triangle\phi|^2 = (h^\phi)^2$; thus bending energy equals Willmore energy of the deformed surface in this case (cf. Section 2).

Linearized elasticity for thin shells While isometric deformations of thin plates lead to quadratic bending energies, this is no longer the case the general case of thin shells, which are not flat in the rest configuration. To gain efficiency by simplifying the elastic energy described above, some authors consider the effect of infinitesimal displacements, i.e., the setting of *linearized elasticity*. In this setting, let $v : \mathcal{S} \to \mathbb{R}^3$ be a vector field on $\mathcal{S}$ with the deformed shape $\phi(\mathcal{S})$ given by

$$\phi(x) = x + \epsilon v(x)$$

for some $\epsilon \in \mathbb{R}$ that is assumed to be small. To further simplify the exposition, we additionally restrict ourselves to energy densities given by the Frobenius norm. i.e., $W_{\text{mem}}(A) = |A - \mathbb{1}|^2$ and $W_{\text{bend}}(A) = |A|^2$. Then the corresponding membrane energy is given by

$$\mathcal{W}_{\text{mem}}[\phi] = \delta\epsilon^2 \int_{\mathcal{S}} |(D_{\mathcal{S}} v)_{\text{sym}}|^2 \, \mathrm{d}a + \dots \,,$$

where dots denote higher order terms with respect to ϵ and $(D_{\mathcal{S}} v)_{\text{sym}}$ is the symmetrized derivative, i.e., $(D_{\mathcal{S}} v)_{\text{sym}} = \mathbb{1}_{3\times 2}^T D_{\mathcal{S}} v + D_{\mathcal{S}} v^T \mathbb{1}_{3\times 2}$, where $\mathbb{1}_{3\times 2}$ is a 3×2 matrix that contains the 2×2 identity matrix in its upper part and zeros in the last row. Likewise, the linearized bending energy for *infinitesimally isometric* (inextensional) deformations is given by

$$\mathcal{W}_{\text{bend}}[\phi] = \delta^3\epsilon^2 \int_{\mathcal{S}} |\text{Hess}[v] \cdot n|^2 \, \mathrm{d}a + \dots \,,$$

where $\text{Hess}(v)$ denotes the component-wise Hessian, and on each component the Hessian is taken with respect to the Riemannian metric on $\mathcal{S}$ (see above). Notice that the inner product with the surface normal n turns $\text{Hess}(v) \cdot n$ into a symmetric 2×2 tensor field. The assumption of infinitesimally isometric reformations is motivated by considering *pure bending*, i.e., by disallowing any contribution of membrane terms in the bending energy.

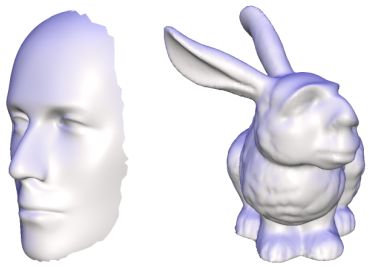

Figure 11: Similar to mesh deformation, the linear editing metaphor can be used for coating transfer (e.g., from mannequin to bunny) [82]. (Image courtesy of O. Sorkine–Hornung.)

Linear elasticity has the benefit of leading to *quadratic* energy functionals and thus *linear* Euler–Lagrange equations *iff* the constraints prescribed by the user (in form of prescribed deformations of parts of the surface) are also linear in positions (cf. Fig. 11). Indeed, in order to be able to work with standard solvers, the above energies are by some authors further modified to

$$\widetilde{\mathcal{W}}_{\text{mem}}[v] = \int_{\mathcal{S}} \|D_{\mathcal{S}} v\|^2 \, \mathrm{d}a \quad \text{and} \quad \widetilde{\mathcal{W}}_{\text{bend}}[v] = \int_{\mathcal{S}} \|\triangle_{\mathcal{S}} v\|^2 \, \mathrm{d}a \,,$$

where $\|\cdot\|$ denotes the usual Euclidean norm. Notice that different from $|(D_{\mathcal{S}} v)_{\text{sym}}|$, which does not account for infinitesimal rotations, $\|D_{\mathcal{S}} v\|$ is the full norm and therefore also incorporates rotations and thus accounts for bending contributions in the membrane energy. Indeed, $\|D_{\mathcal{S}} v\|$ *only* vanishes if v induces a translation. Likewise, taking the trace, i.e., using $\triangle_{\mathcal{S}} v = \text{tr}\,\text{Hess}(v)$, does not account for principal bending directions (see the discussion for parametric surfaces and level sets below). From an implementation point of view, this formulation is simple as it leads to Euler–Lagrange equations that involve standard operators—given by the harmonic and biharmonic equations

$$\triangle_{\mathcal{S}} v = 0 \quad \text{and} \quad \triangle_{\mathcal{S}}^2 v = 0 \,,$$

respectively, for the membrane and bending term. Due to the presence of Laplacian and bi-Laplacian only, this formulation can be used for the variety of surface representations outlined above. We refer the reader to the detailed survey by Botsch and Sorkine–Hornung [14] for a an elaborate discussion of benefits and drawbacks of the linearized approach in computer graphics applications. A major limitation of linear elasticity is that it comes at the price of *loosing rotation invariance*.

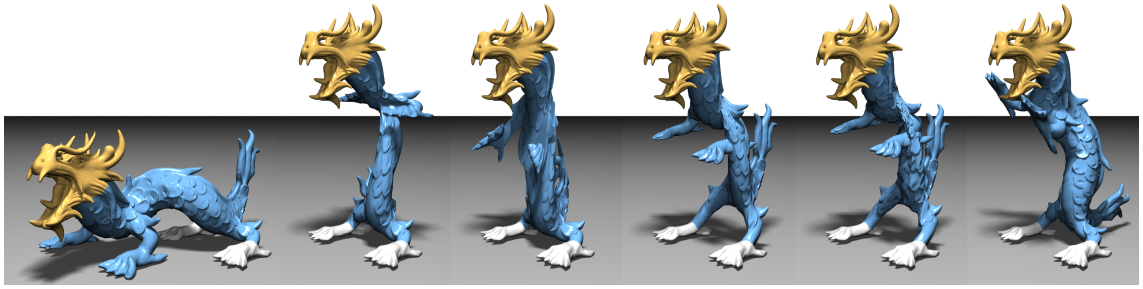

Figure 12: Deformation of the dragon (left) by fixing its hind feet and moving its head upwards in a single step. The result of PriMo [13] (right) compared to linear deformation methods (second through fifth picture) that tend to yield counter-intuitive results. (Image courtesy of M. Botsch.)

Restoring rotation invariance for deforming thin shells Several approaches in the literature have attempted to overcome the undesired limitation of linearized elasticity to violate invariance under rotations.

One set of approaches is based on *moving frames* or what is known as *differential coordinates*. For smooth surfaces, the formulation of representing a surface by moving frames that are adapted to the geometry (i.e., orthonormal frames that have two tangential and one normal unit vector) rather than in a fixed frame (i.e., external Euclidean coordinate system) dates back to E. Cartan and H. Weyl. When using moving frames for deformations, the main idea is to employ a *two-step process*. The methodology of reconstructing a surface in two steps somewhat mimics the proof of the fundamental theorem of surface theory in the smooth setting (when formulated via moving frames), where one also first reconstructs frames (using one set of integrability conditions) and then reconstructs a corresponding surface (using another set of integrability conditions), see [59]. In practice, the first step reconstructs the frame field under the constraints provided by the user. The new frame field is solved for by interpolating (or approximately interpolating) between the frames specified by the user's editing.A second step solves for surface positions using the frame field from the first step and using positional constraints provided by the user. This two-step procedure is not always unproblematic as the constraints of the first and second step may be incompatible. We refer to Botsch and Sorkine–Hornung [14] for a detailed discussion of pros and cons for various methods that employ this methodology.

Another set of approaches builds on giving up linearity to restore rotation invariance. One example of a popular, efficient, and robust method for non-linear surface deformation and shape matching is PriMo [13]. This approach is based on replacing the triangles of a polyhedral surfaces by thick prisms (so as to model a thick surface). During deformation, these prisms are required to stay rigid, while non-linear elastic forces are acting between neighboring prisms to account for bending, twisting, and stretching of the surface. This requires for solving for optimal rigid motions under the constraints prescribed by the user. Unfortunately, there is no immediate continuous analogue of this construction—and it would be interesting to gain more insight, from a continuum mechanics perspective, into a corresponding smooth model. In a similar spirit, Sorkine and Alexa have suggested *as-rigid-as-possible* (ARAP) deformations, which also lacks an evident continuum mechanics analogue [81]. In this approach, one assigns to each vertex an optimal (in a least squares sense) rotation matrix that maps the edges of the undeformed to the edges of the deformed shape. To construct the deformed shape, ARAP alternates between minimizing a suitable energy for optimal vertex positions and optimal rotations.

Volume deformation Another approach that is prevalent in the literature is to view surfaces as boundaries of 3D volumes. In this setting, deformation is handled by changing the entire volume instead of the surface only. From the perspective of continuum mechanics, a suitable energy for isotropic volumetric materials is to minimize the distance of the differential $d\phi$ of a deformation $\phi : M \subset \mathbb{R}^3 \to M' \subset \mathbb{R}^3$ to the rotation group $\mathrm{SO}(3)$, i.e., to minimize the energy

$$\mathcal{W}_{\text{elastic}}[\phi] = \frac{1}{2} \int_M \min_{R(x) \in \mathrm{SO}(3)} |D\phi(x) - R(x)|^2 dx$$

under user-specified constraints. This approach is based on the observation that, apart from *globally* rigid motions, locally varying rotations will unavoidably induce deviation from isometry. Intuitively, this reflects

Figure 13: Snapshots from the simulation of a billowing flag. Despite its economy of cost, the proposed *isometric bending model* in [90] achieves qualitatively the same dynamics as popular nonlinear models.

the fact that if an infinitesimal piece of material is rotated with respect to a "neighboring" infinitesimal piece, this will induce local stretching, compression, or shearing of the material. Notice the conceptual (but unfortunately not precise) similarity of this observation to the approaches employed by PriMo and ARAP discussed above. To be precise, suppose that a volume deformation fixes the center of mass and induces no global rotation, i.e., without loosing generality, by factoring out globally rigid motions, let $\int_M \phi \, \mathrm{d}x = \int_M \operatorname{curl}(\phi) \, \mathrm{d}x = 0$. Then *geometric rigidity* [35] implies that there exists a constant C depending on M (but not on ϕ) such that

$$\int_M |D\phi(x) - \mathbb{1}|^2 \, \mathrm{d}x \leq C \int_M \min_{\mathrm{R(x)} \in \mathrm{SO}(3)} |D\phi(x) - R(x)|^2 \, \mathrm{d}x \, .$$

In comparison to the Green–Lagrange strain tensor $D\phi^T D\phi - \mathbb{1}$, the above formulation is attractive since it is of lower order in the state variables. Chao et al. [19] make use of this observation for constructing an efficient algorithm for volume deformations. We have singled out the above formulation from the bulk of approaches for volume deformations in the literature as it is rotation invariant by design. As for the case of thin shells, rotation invariance would be lost by working with a linearized model.

Physical simulations of thin shells Closely related to surface deformations are physical simulations of shells. There is one important difference, though. While an *eyeball metric* might be a reasonable choice for *deformations* of flexible surfaces for graphics applications, this is often no longer an acceptable metric for *physical simulations*, where, e.g., violation of rotation invariance leads to loss of angular momentum preservation and hence to clearly visible artefacts. The early graphics literature has focused on *efficiency* and *progress*—and, to achieve this, has sometimes been willingly be sacrificing *physical accuracy*. This trend has somewhat been reversed over the past years—with computer science researchers attempting to turn insights from the computational mechanics community into fast algorithms that do not break the laws of physics.

In a physical simulation of thin shells, the state variables ϕ and v, describing the position (as a map from a reference surface $\mathcal{S}$ to $\mathbb{R}^3$) and velocity of a moving surface in space are subject to elastic, damping, contact, and other external forces. The elastic response of a deformed material is governed by a conservative force, i.e., one which acts against the energy gradient:

$$F_{\mathrm{elastic}}[\phi] = -\mathcal{W}'[\phi] = -\mathcal{W}'_{\mathrm{mem}}[\phi] - \mathcal{W}'_{\mathrm{bend}}[\phi] \, . \tag{11}$$

This requires discrete notions of first and second fundamental forms, which we discuss for the prevalent case of polyhedral surfaces.

Damping and collisions Most real materials dissipate energy during motion. *Rayleigh damping* is among the simplest models of dissipation used by the computational mechanics community [48, 93]. In the Rayleigh view, the damping force, F_{damp}, is proportional to velocity, $v(t) = \dot{\phi}(t)$:

$$F_{\mathrm{damp}}[v] = -K_{\mathrm{d}} v \, , \quad \text{where} \quad K_{\mathrm{d}} = \alpha_1 M + \alpha_2 \mathrm{Hess} \, . \tag{12}$$

The constant of proportionality is written as a linear combination[1] of two tensors: the mass matrix and the Hessian of the elastic energy; the two tensors correspond to damping of low and high temporal frequencies, respectively.

A far more challenging and involved problem for physical simulations is *collisions*. Indeed, while the problem of how to *detect* collisions of a surface (either with other objects or with parts of itself) has been treated by efficient algorithms in the literature (e.g., by using specifically adapted spatial data structures), the problem of how to *resolve* collisions in a physically accurate manner is still an active area of research. Consider scenarios such as multiple contact regions that frequently change over time or sliding contact, e.g., when pulling a rope tight. Challenging aspects include that colliding regions must not interpenetrate, that collisions are to be resolved in physically correct way, and that computations for resolving collisions should finish in finite time. For a fairly recent treatment of these aspects we refer to the work of Harmon et al. [40].

Temporal evolution In a classical mechanical system, the temporal evolution of position, $\phi \equiv \phi(t)$, and velocity, $v \equiv \dot{\phi}$, is governed by the equations of motion:

$$\begin{pmatrix} \dot{\phi}(t) \\ \dot{v}(t) \end{pmatrix} = \begin{pmatrix} Id & 0 \\ 0 & (\rho M)^{-1} \end{pmatrix} \begin{pmatrix} v(t) \\ F_{\text{elastic}}[\phi(t)] + F_{\text{damp}}[v(t)] + F[\phi(t), v(t))] \end{pmatrix}, \tag{13}$$

with initial conditions $\phi(0)$ and $v(0)$. Here $F[\phi, v]$ denotes other forces, such as contact forces or gravity, and the *physical mass matrix*, ρM, is given by the product of mass surface density and the geometric mass matrix. In order to accelerate force computation, it is often desirable to have *explicit* representations of energy gradients and energy Hessains instead of numerically deriving these quantities on the fly by automatic differentiation.

Time discretization of (13) is a well-studied problem (see [41] and references therein); approaches may be classified as explicit, implicit, or mixed implicit-explicit.[2] *Geometric (or variational) integrators* [57] have been advocated due to their *structure preservation*, i.e., guaranteed preservation of momenta and near preservation of energy even for large times.

6 Matching of Thin Shells

The feature aware matching of two given surfaces is one of the fundamental tasks is geometry processing. The matching problem consists of finding a "good" correspondence between two given input shapes, such as two faces, two poses of an animated character, or two scans of body organs. The meaning of "good" is dependent on the specific application and is often measured with the help of some energy functional. One challenge in shape matching is that in principle *all* possible correspondences between the two given shapes would have to be considered—a space that is often intractably large for computations. Another challenge for shape matching based on energy minimization is the nonconvexity of the attendant energy landscape—a problem that can be alleviated using multiscale or multigrid optimizers. Finally, a big challenge is *partial shape matching* in scenarios where shapes are only partially available (e.g., matching a humanoid figure to a hand). There exists a vast amount of literature on matching triangulated surfaces. For example, starting from the notion of Gromov–Hausdorff distances Bronstein et al. [16] have developed an efficient algorithm denoted as *multidimensional scaling* for the matching of triangular surfaces also for the case of partial correspondence. A topologically robust variant of this approach based on diffusion distances is proposed in [15]. The observation that robustness to noise and changes of topology can be significantly increased by a multiscale approach has been utilized by Sun, Ovsjanikov and Guibas in [84], where shape matching is based on the heat kernel.

Here we investigate the shape matching problem from the viewpoint of thin shell deformations and under two different perspectives: the matching of parametric surfaces using deformations between the parameter domains and the matching of implicit surfaces.

Matching parametric surfaces. If parametrizations of the surfaces are given,

[1] In this ad-hoc model, the constants α_1 and α_2 are endowed with the requisite units so that the final product has units of force.

[2] In mixed implicit-explicit (IMEX) time-integration, some forces are treated using the explicit method, and other forces are treated using the implicit method.

the problem can be phrased in terms of a functional on matching deformations ϕ^ω from the parameter domain ω_A of the surface $\mathcal{S}_A$ to $\mathbb{R}^2$ where only partial matching ($\phi(\omega_A) \neq \omega_B$) is allowed with respect to the parameter domain ω_B of the surface $\mathcal{S}_B$. Physically, one might consider this as the pressing of a given surface $\mathcal{S}_A$ into the mould of the surface $\mathcal{S}_B$. Besides a tangential distortion energy $\mathcal{W}_{\text{mem}}$, which can be regarded as a regularization energy ensuring smoothness of the matching deformation, and a bending energy $\mathcal{W}_{\text{bend}}$, which tries to match regions of equal curviness such as corresponding creases, one finally might want to align preselected feature regions like the eyes on facial surfaces via a third energy term $\mathcal{W}_{\text{feature}}$. In detail, the tangential distortion energy can be chosen as

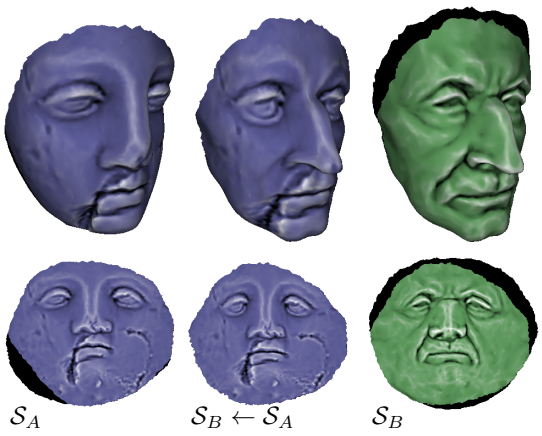

Figure 14: In surface matching, a partial correspondence is often desired. The correspondence is defined where their parameter domains intersect under the matching deformation (bottom). In this domain, quantities such as the mean curvature texture maps can be mapped between the surfaces (center). The unmatched regions are in black.

$$\mathcal{W}_{\text{mem}}[\phi^\omega] = \int_{\omega_A} W_{\text{mem}}(\mathbf{G}^\omega[\phi^\omega]) \sqrt{\det g_A}\, d\xi\,.$$

where $W_{\text{mem}}(A) = \hat{W}(a,d) = \alpha_l a + \alpha_a\big(d + (1+\frac{\alpha_l}{\alpha_a})\, d^{-1}\big)$. Here, $a = a(A) = \operatorname{tr} \mathbf{G}^\omega[\phi]$ accounts for length distortion and $d = d(A) = \det \mathbf{G}^\omega[\phi]$ for area expansion with and area compression with d^{-1}. The weights $\alpha_l, \alpha_a > 0$ are chosen according to the relative importance of length and area distortion. This simple class of polyconvex energy functionals [25] was rigorously derived in [60] from a set of natural axioms for measuring the distortion of a single parameterization. Bending of the surface $\mathcal{S}_A$ under the deformation $\phi = x_B \circ \phi^\omega \circ x_A^{-1}$ is penalized by the bending energy (10). We obtain for the relative shape operator

$$\mathbf{S}^\omega_{\text{rel}} = g_A^{-1}(D(\phi \circ x_A)^T D(n^\phi \circ \phi \circ x_A) - Dx_A^T D(n \circ x_A)) = g_A^{-1}(D\phi^\omega)^T(\mathbf{S}^\omega_B \circ \phi^\omega)\, D\phi^\omega - \mathbf{S}^\omega_A\,,$$

and the resulting functional on the deformation ϕ^ω in a general form

$$\mathcal{W}_{\text{bend}}[\phi^\omega] = \int_{\omega_A} \chi_B \circ \phi^\omega F(\mathbf{S}^\omega_{\text{rel}}) \sqrt{\det g_A}\, d\xi$$

for some function $F : \mathbb{R}^{2,2} \to \mathbb{R}$ and the characteristic function χ_B of $x_B(\omega_B)$. If one is only interested in the comparison of the mean curvatures of the surfaces defined as the traces of the shape operators, a simplified version of the functional is given by $\mathcal{W}_{\text{mem}}]\phi^\omega] = \int_{\omega_A} \chi_B \circ \phi^\omega |h_A - h_B \circ \phi^\omega|^2 \sqrt{\det g_A}\, d\xi$ with $h_A = \operatorname{tr}(\mathbf{S}_A)$, $h_B = \operatorname{tr}(\mathbf{S}_B)$. Finally, if we denote by $\mathcal{F}_A \subset \omega_A$ and $\mathcal{F}_B \subset \omega_B$ features sets in the parameter domains of the two surfaces, then

$$\mathcal{W}_{\text{feature}}[\phi^\omega] = \int_{\omega_A} ((\chi_{\mathcal{F}_B} \circ \phi^\omega)(1 - \chi_{\mathcal{F}_A}) + (1 - \chi_{\mathcal{F}_B} \circ \phi^\omega)\chi_{\mathcal{F}_A})\, \sqrt{\det g_A}\, d\xi$$

measures the symmetric difference on the surface A of the feature set on surface A and the pull back of the feature set on the surface B. The resulting combined energy is then given by $\mathcal{W}[\phi^\omega] = \alpha_{\text{mem}} \delta \mathcal{W}_{\text{mem}}[\phi^\omega] + \alpha_{\text{bend}} \delta^3 \mathcal{W}_{\text{bend}}[\phi^\omega] + \alpha_{\text{feature}} \mathcal{W}_{\text{feature}}[\phi^\omega]$ with suitable weights. Fig. 14 shows a blending application of surface matching.

Matching implicit surfaces. Frequently, surfaces are described and modeled as level set of a function w on a computational domain $D \subset \mathbb{R}^3$. For fixed surfaces $\mathcal{S}_A$ and $\mathcal{S}_B$ it is advantages to represent them by signed distance functions d_A and d_B, respectively. In particular in this case one obtains $\mathbf{S}_X = Dn_X P_X = D^2 d_X$ for $P = \mathbb{1} - \nabla d_X \otimes \nabla d_X$ and $X \in \{A, B\}$. For the conversion of a general implicit surface representation into a signed distance function a robust and efficient algorithm is the fast marching method [80]. If one is just interested in the matching of the two surfaces $\mathcal{S}_A$ and $\mathcal{S}_B$ a narrow band approach enables

to "blend out" distant implicit surfaces $[d_A = c]$ and $[d_B = c]$ for $|c| \geq \epsilon$. In fact, one defines a smooth blending function $\eta_\sigma : \mathbb{R} \to \mathbb{R}_0^+$ with $\eta_\sigma(c) = 0$ for $|c| \geq \sigma$ and $\eta_\sigma(c) = 1$ for $|c| \leq \frac{\sigma}{2}$. Then, taking into account the coarea formula ($\nabla d_A = 1$) one obtains for a deformation $\phi : D \to D$ the tangential distortion energy on the narrow band $[\sigma \leq d_A \leq \sigma]$

$$\mathcal{W}_{\mathrm{mem}}[\phi] = \int_D (\eta_\sigma \circ d_A)\, W_{\mathrm{mem}}(\mathbf{G}[\phi] + \nabla d_A \otimes \nabla d_A)\, \mathrm{d}x\,,$$

where $\mathbf{G}^{T_x\mathcal{S}}[\phi]$ and W_{mem} are the above defined tangential Cauchy Green strain tensor and the associated energy density, respectively. Here incorporated projection operator P is the one corresponding to the distance function d_A, i.e., $P = P_A$. As a first approach for the bending energy one gets [51]

$$\mathcal{W}_{\mathrm{bend}}[\phi] = \int_D \eta_\sigma(d_A) \| P_A D\phi^T ((D^2 \mathrm{dist}_B) \circ \phi)\, D\phi P_A - D^2 d_A \|^2\, \mathrm{d}x\,.$$

Unfortunatly, the resulting energy is not quasi convex and hence not lower semicontinuous. To obtain a lower semicontinuous energy, we modify the definition of the tangential Jacobian setting $D_{AB}\phi = (P_B \circ \phi) D\phi P_A$ which coincides with the standard definition on the surface $\mathcal{S}_A$ if $\phi(\mathcal{S}_A) \subset \mathcal{S}_B$. Then, one obtains the correspondingly modified Cauchy Green strain tensor $\mathbf{G}_{AB}^{T_x\mathcal{S}}[\phi] = D_{AB}\phi^T D_{AB}\phi$. Now, let us assume that $\underline{\beta}$ is a strict lower bound for the negative principal curvatures on all involved surfaces and 0 if the principal curvatures are all positive. Then we define

$$\mathcal{W}_{\mathrm{mem}}[\phi] = \int_D (\eta_\sigma \circ d_A) W_{\mathrm{bend}} \left(((S_B - \underline{\beta}\mathbb{1})^{\frac{1}{2}} \circ \phi)\, D_{AB}\phi\, (S_A - \underline{\beta}\mathbb{1})^{-\frac{1}{2}} \right) \mathrm{d}x\,,$$

which now is indeed a lower semicontinuous functional [50]. Here, the integrant W_{bend} is chosen similar to the integrant of the membrane energy. Besides the membrane and bending energies on the narrow band, one needs a penalty functional

$$\mathcal{W}_{\mathrm{penalty}}[\phi] = \frac{1}{\epsilon} \int_D \eta_\sigma(D_A)(\mathrm{dist}_A - \mathrm{dist}_B \circ \phi)^2\, \mathrm{d}x$$

to ensure that the neighboring level sets around $\mathcal{S}_A$ in the narrow band are sufficiently well matched to the neighboring level sets of surface $\mathcal{S}_B$. Finally, to render the method computationally feasible one has in addition to ensure that the deformation ϕ is regular and injective also outside the narrow band. To this end, we incorporate a regularization energy

$$\mathcal{W}_{\mathrm{reg}}[\phi] = \int_D (1 - \eta_\sigma(D_A)) W_{\mathrm{reg}}(C[\phi])\, \mathrm{d}x$$

where $C[\phi] = D\phi(x)^T\, D\phi(x)$ is the three-dimensional Cauchy-Green strain tensor and $W_{\mathrm{reg}}(C) = \frac{\mu}{4}\mathrm{tr}(C)^2 + (\lambda - 2\mu)\det(C)^{\frac{1}{2}} - \frac{\lambda-\mu}{2}\log\det(C)$ for Lamé-Navier coefficients λ and μ [20]. Thus, combining the different energy terms we finally obtain the total energy

$$\mathcal{W}[\phi] = \alpha_{\mathrm{mem}}\delta\mathcal{W}_{\mathrm{mem}}[\phi] + \alpha_{\mathrm{bend}}\delta^3\mathcal{W}_{\mathrm{bend}}[\phi] + \alpha_{\mathrm{penalty}}\mathcal{W}_{\mathrm{penalty}}[\phi] + \alpha_{\mathrm{reg}}\mathcal{W}_{\mathrm{reg}}[\phi]\,.$$

Like in the the parametric case, a feature matching energy could additionally be taken into account.

7 Spectral and Modal Methods

Building on a general paradigm of global analysis one studies the spectrum of differential operators on surfaces and considers the spectrum and the eigenfunctions as descriptors of the surface. A prominent operator is again the Laplace–Beltrami operator that allows for applying classical approaches to spectral analysis from signal processing to irregular triangle meshes. The eigenfuctions of the Laplacian constitute a basis that provides the analogue of the Fourier basis in the planar case. Spectral geometry processing applications include mesh compression [8, 52], parameterization [67], mesh fingerprints [76], mesh segmentation and registration [75], and so forth, to name a few.

Another operator that is relevant from the perspective of shell deformations is the *Hessian* of the elastic energy discussed in Section 6. If the Hessian of elastic energy is considered at the (undeformed) rest state

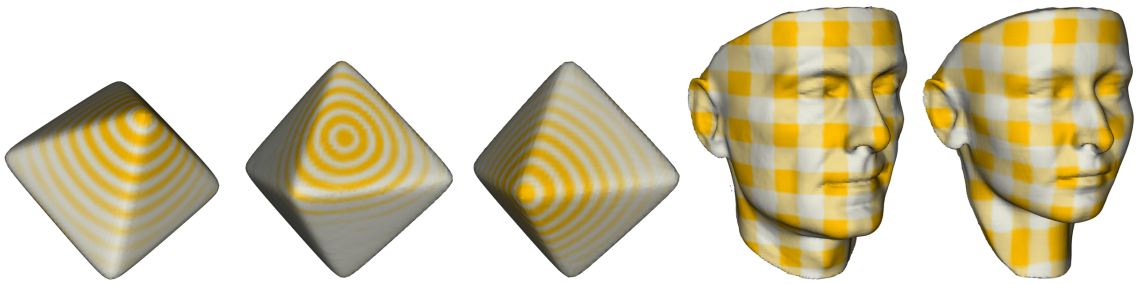

Figure 15: Matching of implicit surfaces based on the simplified thin shell model [51] for two octahedrons (left) (the second image shows a matching without bending energy) and two different faces (right). Thereby the texture is deformed together with the template surface and shows the proper matching of geometric features.

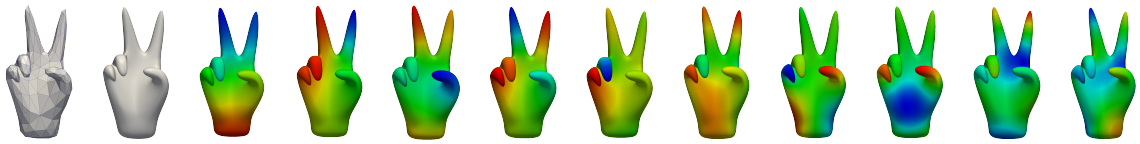

Figure 16: The first 10 eigenfunctions are color coded on the limit surface of Loop's subdivision applied to an initial control mesh (left) (Image courtesy of R. Perl).

of a (possibly naturally curved) shell, then the eigenvalues λ_i are eigenfrequencies and the corresponding eigenfunctions $u_i : \mathcal{S} \to \mathbb{R}^3$ are the associated *vibration modes.* Thereby a vibration mode u_i is a displacement function on the surface and describes an (infinitesimal) oscillation of the elastic shell with frequency λ_i. Thus, a low eigenvalue corresponds to a low degree of stiffness in that direction and thus a physically preferred mode of shell variability.

In general the Hessian of a function requires the notion of a Riemannian metric, i.e., for a function $f : \mathcal{M} \to \mathbb{R}$ on a Riemannian manifold$(\mathcal{M}, g)$ the Hessian $\mathrm{Hess}(f)$, evaluated at a point $p \in \mathcal{M}$ requires the notion of a Riemannian metric g and *changes* if g is changed. However, if the differential of f vanishes at p, then the Hessian turns out to be *independent* of the choice of metric. This observation implies that it is meaningful to consider the Hessian of elastic energy *when evaluated at the rest state*. Indeed, let $\mathcal{S} \subset \mathbb{R}^3$ be the undeformed surface describing a thin shell, and consider an energy $\mathcal{W}$ of surface deformations $\phi : \mathcal{S} \to \mathbb{R}^3$. Suppose that the identical deformation $\mathbb{1}$ of $\mathcal{S}$ is a critical point of the energy and therefore $\mathcal{W}'[\mathbb{1}](v) = 0$ for all displacement fields $v : \mathcal{S} \to \mathbb{R}^3$, which holds true for the shell energies derived in Section 6. Regarding a surface $\mathcal{S}$ as a *point* in shape space, observe that a displacement field $v : \mathcal{S} \to \mathbb{R}^3$ can be considered as *tangent vector* at $\mathcal{S}$ in the space of shapes. Due to the criticality of $\mathcal{W}$ at the point $\mathcal{S}$ it follows that the Hessian of $\mathcal{W}$ satisfies $\mathrm{Hess}\,\mathcal{W}[\mathbb{1}](v, w) = \mathcal{W}''[\mathbb{1}](v, w)$, *independently* of any choice of Riemannian structure on the space of shapes. Thus, independent of the choice of a Riemannian metric on the space, $\mathrm{Hess}\,\mathcal{W}[\mathbb{1}](v, w)$ is a symmetric endomorphism on the space of displacement fields representing the second variation of the energy.

The spectrum and eigenvalues of this Hessian provide the so-called *modal basis*, which is in general different from a spectral basis, and that is adapted to elastic propertied of a given surface. Hildebrandt et al. [46, 89] showed how to use this basis for the intuitive modeling of surfaces and for accelerating physical simulations.

As a significantly simplified, linear model one considers solely normal variations $u\,n$ of the surface $\mathcal{S}$ and the quadratic energy $\mathcal{W}[u] = \int_{\mathcal{S}} |\triangle_{\mathcal{S}} u|^2 \,\mathrm{d}a$ for scalar functions $u : \mathcal{S} \to \mathbb{R}$. Then, one is lead to the eigenvalue problem $\triangle_{\mathcal{S}}^2 u = \lambda u$ for the geometric bi-Laplacian $\triangle_{\mathcal{S}}^2$. A set of eigenfunctions for the geometric bi-Laplacian is shown in Figure 16. Here, a projected inverse vector iteration has been applied to the discrete weak form (6) in the Loop subdivision finite element approach.

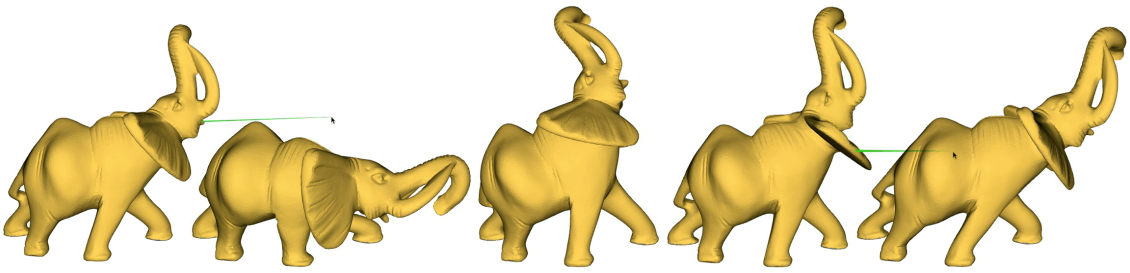

Figure 17: Results of an interactive animation of an elephant model using modal analysis to build a reduced basis of forces [89] (Image courtesy of K. Hildebrandt).

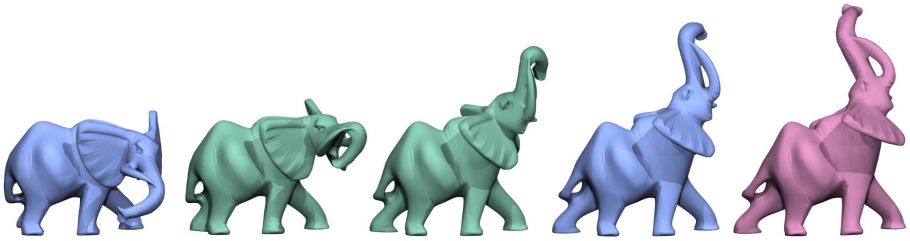

Figure 18: Two intermediate surface (green) from a geodesic between two poses (blue) of an elephant model and a extrapolation via the exponential map (magenta) using the approach by Kilian *et al.* [55] (Image courtesy of M. Kilian)

8 From shapes to the space of shapes

So far, we have discussed computational tools to process single geometries. From a more global perspective one might want to study surfaces as objects in a space of shapes. In particular over the last decade, concepts from Riemannian manifolds have been applied to design and investigate nonlinear and frequently infinite-dimensional shape spaces, with applications in shape morphing and modeling [55, 56], in computational anatomy [6, 65], as well as shape statistics [33, 36].

For planar *curves*, different Riemannian metrics have been devised, including curvature-weighted L^2- or Sobolev-type metrics [64, 86]. Some of these metrics can be phrased as measures of stretching and bending [83, 85]. In a geometrically motivated approach, Kilian *et al.* consider geodesics between consistently triangulated surfaces [55], where the underlying Riemannian metric measures the stretching of triangle edges (cf. Fig. 18). The resulting metric is intrinsically invariant with respect to isometric deformations; however, due to the lack of a bending energy and in order to select a smooth path in the set of different isometric paths, a supplementary regularization is required. Figure 18 shows surfaces from a corresponding geodesic path.

Building on the work of Kilian et al. [55], we investigate the space of thin shell surfaces as a Riemannian shape manifold $\mathcal{M}$. The central observation is that the Hessian of elastic energy described in Section 7, when evaluated at the undeformed surface is *positive semidefinite* and its kernel consists of infinitesimal rigid transformations only. Thus, modulo rigid transformations of surfaces, this Hessian provides a notion of a Riemannian metric on the space of shapes. At a pint $\mathcal{S}$ in shape space, we denote this metric by $\mathbf{g}_{\mathcal{S}}(\cdot,\cdot)$.

The geometric structure of shape space is described in terms of a Riemannian metric $\mathbf{g}_{\mathcal{S}}(\cdot,\cdot)$, which assigns a cost to any inifinitesimal shape variations and thus encodes directions of preferred shape variability. Riemannian shape calculus provides a set of basic tools useful in applications: the computation of shortest geodesic paths $(\mathcal{S}(t))_{t\in[0,1]}$ between surfaces $\mathcal{S}_A = \mathcal{S}(0)$ and $\mathcal{S}_B = \mathcal{S}(1)$ as minimizers of the path energy $\int_0^1 \mathbf{g}_{\mathcal{S}(t)}(\dot{\mathcal{S}}(t),\dot{\mathcal{S}}(t))\,\mathrm{d}t$ establishes a rigorous notion of distance between surfaces; the logarithmic map allows to represent large, nonlinear shape variability in terms of an infinitesimal shape variation in the tangent space of the shape manifold; the exponential map allows geometrically and physically sound shape modeling and animation starting from a infinitesimal variation of a given shape; and finally geometric details can be appropriately transferred from one surface onto another surface using parallel transport.

When using a Riemannian metric based on the Hessian of elastic energy, it turns out that the attendant

Figure 19: Discrete 8-geodesic between two hand poses, based on the model of discrete viscous shells. The colors indicate dissipation due to surface stretching (top) and bending (bottom).

geodesics have a physical interpretation. Indeed, wiewing surfaces as thin *viscous* materials of some finite (albeit small) thickness, a geodesic is the deformation path between two points in the space of shells that has the *least energy dissipation*, i.e., that requires the least total effort to undergo the deformation. Let us view shells as a thin *three-dimensional* objects of thickness δ with $\mathcal{S}$ describing the midsurface, whereas $\mathcal{S}^\delta$ denotes the corresponding material object around $\mathcal{S}$. Now, A family $(\phi_t)_{0\leq t\leq 1}$ of diffeomorphisms with $\phi_0 = \mathrm{Id}$ generates a deformation path $\phi_t(\mathcal{S}^\delta)$. If we assume that the shell were made of *viscous* material, then applying such a family of deformation leads to viscous friction within the shell's volume. The density of this friction—known as dissipation—is given by a *quadratic* function $\mathrm{Diss}(\epsilon[v])$ of the symmetrized gradient $\epsilon[v] := \frac{1}{2}\left((\mathrm{D}v)^T + \mathrm{D}v\right)$ of the associated Eulerian velocity field $v(t) = \dot{\phi}(t) \circ \phi(t)^{-1}$. We treat viscous dissipation according to Rayleigh's analogy that derives a viscous formulation from an elastic one by replacing elastic strain by strain *rates* [74]. To this end, we retrieve the *elastic* energy density W acting on symmetric strain tensors $A \in \mathbb{R}^{3,3}$. We assume that $A = \mathbb{1}$ is a minimizer of W, which implies that $W(\mathbb{1}) = 0$ and $W_{,A}(\mathbb{1}) = 0$. According to Rayleigh's analogy, the dissipation density is now given by the second derivative of elastic energy density at the identity, i.e., $\mathrm{Diss}(C) = \frac{1}{2}W_{,AA}(\mathbb{1})(C, C)$ for any field of symmetric tensors $C \in \mathbb{R}^{3,3}$. Here, we view $C = \epsilon[v]$ as the description of a rigid body motion invariant infinitesimal variation of the shell object $\mathcal{S}^\delta$. Thus, we choose as a Riemannian metric the spatially integrated dissipation density $\mathbf{g}_{\mathcal{S}^\delta}(C, C) = \mathrm{Diss}(C)$. The the resulting path energy—total energy dissipation associated with the *accumulated* viscous friction along the deformation path—is then given by

$$\mathcal{E}[(\phi(t))_{0\leq t\leq 1}] = \int_0^1 \int_{\phi(t,\mathcal{S}^\delta)} \mathbf{g}_{\mathcal{S}^\delta}(\epsilon[v](t,x), \epsilon[v](t,x))\,\mathrm{d}x\,\mathrm{d}t\,. \tag{14}$$

Figure 19 shows a discrete geodesic together with the different component of the underlying local dissipation rate.

Notwithstanding the conceptual power of a Riemannian calculus in shape spaces, operators such as the logarithm or the exponential map involve the solution of time-dependent nonlinear ordinary or partial differential equations. In more complex spaces—as they appear in vision applications—solutions to these equations are typically difficult to compute. Here, a discrete theory provides a corresponding set of time discrete tools comparatively simple to state and to implement. It is centered around the definition of a discrete geodesic as the minimizer of a time-discrete path energy and naturally extends from there to discrete analogs of logarithm, exponential map, and parallel transport. Here, we give a very brief outline and motivate the concept. For details we refer to [77, 78]. In order to transform the above notion of time-continuous geodesics into a corresponding time-discrete one, consider the time-discrete family $(\mathcal{S}_k^\delta)_{k=0,\dots,K}$ of shells given by $\mathcal{S}_k^\delta = \phi_{k\tau}(\mathcal{S}^\delta)$, where $\tau = \frac{1}{K}$ denotes the discrete time step. According to the above discussion of Rayleigh's analogy, we consider the *total elastic energy* $\mathcal{W}[\mathcal{S}^\delta, \phi(\mathcal{S}^\delta)] = \int_{\mathcal{S}^\delta} W(\mathrm{D}\phi^T \mathrm{D}\phi)\,\mathrm{d}x$ associated with a diffeomorphic deformation ϕ from $\mathcal{S}^\delta$, where $\mathrm{D}\phi^T\mathrm{D}\phi$ is the Cauchy Green strain tensor. Now, a straightforward Taylor expansion shows that one can approximate, up to second order in τ, the path energy

(14) by

$$\mathcal{E}[\mathcal{S}_0,\dots,\mathcal{S}_K] = K\sum_{k=1}^{K}\mathcal{W}[\mathcal{S}_{k-1}^{\delta},\phi_k(\mathcal{S}_{k-1}^{\delta})]\,, \tag{15}$$

where $\phi_k := \phi_{k\tau}\circ\phi_{(k-1)\tau}^{-1}$ and $\phi_k(\mathcal{S}_{k-1}^{\delta}) = \mathcal{S}_k^{\delta}$. Notice that on the left-hand side of (15) we have tacitly represented the corresponding material volume $\mathcal{S}_k^{\delta}$ by the corresponding midsurfaces $\mathcal{S}_k$. A discrete geodesic path is now defined as a minimizes of the discrete path energy $\mathcal{E}[\mathcal{S}_0,\dots,\mathcal{S}_K]$ with $\mathcal{S}_0 = \mathcal{S}_A$ and $\mathcal{S}_K = \mathcal{S}_B$ are given. Finally, combining membrane and bending contributions of the midsurace, we replace the deformation energy by

$$\mathcal{W}[\mathcal{S},\phi(\mathcal{S})] = \left(\delta\mathcal{W}_{\mathrm{mem}} + \delta^3\mathcal{W}_{\mathrm{bend}}\right)[\mathcal{S},\phi(\mathcal{S})] \tag{16}$$

and rewrite the path energy correspondingly taking into account the two dimensional shell surfaces. Concerning the spatial discretization for triangulated meshes, one might use the notions of discrete first and second fundamental forms discussed in Section 2.

In what follows, let us denote the vector of vertex positions uniquely describing a discrete surface $\mathcal{S}_h$ by $\bar{X} = \bar{X}[\mathcal{S}_h]$. Thus, we describe discrete surfaces solely in terms of these vectors and adapt the notion of the spatially discrete deformation energy and the fully discrete path energy accordingly, i.e., using the notion $\mathcal{W}^h[\bar{X},\bar{Y}]$ and $\mathcal{E}^h[\bar{X}_0,\dots,\bar{X}_K]$, respectively. If $(\bar{X}_0,\dots,\bar{X}_K)$ is a discrete geodesic, then the initial displacement $\bar{X}_1 - \bar{X}_0$ scaled with $\frac{1}{\tau}$ can be considered as a discrete counterpart of the continuous geometric logarithm $\log_{\bar{X}_A}\bar{X}_B$, which gives rise to the notation $\frac{1}{K}\mathrm{Log}_{\bar{X}_A}\bar{X}_B = \bar{X}_1 - \bar{X}_0$. The continuous geometric exponential map $\exp_{\bar{X}}v$ is defined as the endpoint of a geodesic with initial velocity v in the tangent space at the position $\bar{X}$ on the shape manifold. Thus, for a given displacement $\bar{V}$ one defines a discrete exponential map $\mathrm{Exp}_{\bar{X}}^{K}\bar{Y}$ (with K time steps) as that position $\bar{Y}$, such that the initial displacement of the resulting discrete geodesic $[\bar{X} = \bar{X}_0,\dots,\bar{X}_K = \bar{Y}]$ between $\bar{X}$ and $\bar{Y}$ equals $\bar{V}$, i.e., $\bar{V} = \bar{X}_1 - \bar{X}_0$. Figure 20 shows results of the discrete exponential map. Finally, let us perspectively outline how to proceed with respect to a discrete parallel transport. To define a discrete parallel transport $\mathbf{P}_{\bar{X}_0,\dots,\bar{X}_K}\bar{V}$ of a displacement $\bar{V}$ along a discrete curve $(\bar{X}_0,\dots,\bar{X}_K)$ one might use a discrete counterpart of a Riemannian parallelogram construction. In fact, one proceeds iteratively along the discrete path, starting with the displacement $\bar{V}_0 = \bar{V}$ at the position $\bar{X}_0$. Then, for a displacement $\bar{V}_{k-1}$ at position $\bar{X}_{k-1}$ the transported displacement $\bar{V}_k$ at position $\bar{X}_k$ can be computed as follows. At first, a discrete geodesic $(\bar{X}_{k-1}+\bar{V}_{k-1},\bar{X}_k^c,\bar{X}_k)$ with a midpoint position $\bar{X}_k^c$ is computed. Then, one computes $\bar{X}_k^+ = \mathrm{Exp}^2_{\bar{X}_{k-1}}\bar{X}_k^c - \bar{X}_{k-1}$ to obtain a discrete geodesic with midpoint $\bar{X}_k^c$ starting at the position $\bar{X}_{k-1}$. Finally, one evaluates $\bar{V}_k = \bar{X}_k^+ - \bar{X}_k$. Indeed, the resulting two discrete geodesics will represent discrete diagonals of a Riemannian parallelogram with vertex positions $\bar{X}_{k-1}$, $\bar{X}_k$, $\bar{X}_k+\bar{V}_k$ and $\bar{X}_{k-1}+\bar{V}_{k-1}$.

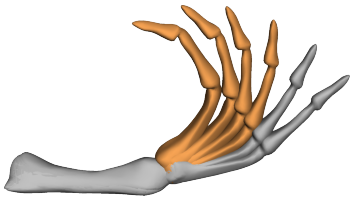

Figure 20: Two positions of a finger as initial shapes (grey) and Superimposed several steps of discrete exponential map.

9 Conclusions

In this paper we have highlighted some of the developments in geometry processing with a strong focus on methods related to the mechanics of thin elastic surfaces and from a naturally strongly biased personal perspective. We tried to cover various geometric representations including parametric and triangulated surfaces, point clouds, level sets, and subdivision surfaces. Furthermore, we have sketched the increasing interaction of mathematics, computer graphics, and mechanics. This interplay is characterized by fundamental and mainly unresolved challenges. Indeed, the convergence of discrete minimizers of geometric and elastic functionals to their continuous counterparts is widely open. Likewise it is unknown for the elastic functionals describing surface deformations if their exist discrete (local) minimizers in the vicinity of a continuous (local) minimizer. Furthermore, concerning the dynamics of shapes the aim is to advance consistent physical animation towards fully fletched virtual reality. Finally, while shapes are quite well understood enrolling the structure of the space of shapes posed many so far non tackled questions.

References

[1] M. Alexa and M. Wardetzky. Discrete laplacians on general polygonal meshes. *ACM Transaction on Graphics*, 30:102:1–102:10, 2011.

[2] N. Amenta, S. Choi, T. Dey, and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. In *ACM Sympos. Comput. Geom.*, pages 213–222, 2000.

[3] D. N. Arnold, R. S. Falk, and R. Winther. Finite element exterior calculus, homological techniques, and applications. *Acta Numerica*, 15:1–155, 2006.

[4] C. Bajaj and G. Xu. Anisotropic diffusion of subdivision surfaces and functions on surfaces. *ACM Transactions on Graphics*, 22(1):4–32, 2003.

[5] T. Baraff and A. Witkin. Large steps in cloth simulation. In *Proceedings of ACM SIGGRAPH*, pages 43–54, 1998.

[6] M. F. Beg, M. Miller, A. Trouvé, and L. Younes. Computational anatomy: Computing metrics on anatomical shapes. In *Proceedings of 2002 IEEE ISBI*, pages 341–344, 2002.

[7] M. Belkin, J. Sun, and Y. Wang. Constructing Laplace operator from point clouds in r^d. SODA '09, pages 1031–1040, 2009.

[8] M. Ben-Chen and C. Gotsman. On the optimality of spectral compression of mesh data. *ACM Trans. Graph.*, 24(1):60–80, 2005.

[9] A. I. Bobenko and R. Seiler. *Discrete Integrable Geometry and Physics*. Oxford University Press, 1999.

[10] A. I. Bobenko, J. M. Sullivan, P. Schröder, and G. Ziegler. *Discrete Differential Geometry*. Oberwolfach Seminars. Birkhäuser Basel, 2008.

[11] A. I. Bobenko and Y. B. Suris. *Discrete Differential Geometry. Integrable Structure*. American Mathematical Society, 2008.

[12] J.-D. Boissonnat. Geometric structures for three-dimensional shape representation. *ACM Transactions on Graphics*, 3:266–286, 1984.

[13] M. Botsch, M. Pauly, M. Gross, and L. Kobbelt. PriMo: Coupled prisms for intuitive surface modeling. In *Proceedings of SGP*, pages 11–20, 2006.

[14] M. Botsch and O. Sorkine. On linear variational surface deformation methods. *IEEE Transactions on Visualization and Computer Graphics*, 14:213–230, 2008.

[15] A. M. Bronstein, M. M. Bronstein, R. Kimmel, M. Mahmoudi, and G. Sapiro. A Gromov-Hausdorff framework for topologically-robust non-rigid shape matching. *IJCV*, 89:266–286, 2010.

[16] M. Bronstein, M. M. Bronstein, and R. Kimmel. Generalized multidimensional scaling: A framework for isometry-invariant partial surface matching. *PNAS*, 103:1168–1172, 2006.

[17] E. Catmull and J. Clark. Recursively generated b-spline surfaces on arbitrary topological meshes. *Computer Aided Design*, 10:350–355, 1978.

[18] F. Catté, P.-L. Lions, J.-M. Morel, and T. Coll. Image selective smoothing and edge detection by nonlinear diffusion. *SIAM J. Numer. Anal.*, 29(1):182–193, 1992.

[19] I. Chao, U. Pinkall, P. Sanan, and P. Schröder. A simple geometric model for elastic deformations. In *Proceedings of ACM SIGGRAPH*, pages 38:1–38:6, 2010.

[20] P. G. Ciarlet. *Mathematical Elasticity, Volume I: Three-dimensional elasticity*, volume 20 of *Studies in Mathematics and its Applications*. Elsevier, 1988.

[21] F. Cirak, M. Ortiz, and P. Schröder. Subdivision surfaces: a new paradigm for thin-shell finite-element analysis. *International Journal for Numerical Methods in Engineering*, 47(12):2039–72, 2000.

[22] D. Cohen-Steiner and J.-M. Morvan. Restricted Delaunay triangulations and normal cycle. *Sympos. Comput. Geom.*, pages 312–321, 2003.

[23] D. Cohen-Steiner and J.-M. Morvan. Second fundamental measure of geometric sets and local approximation of curvatures. *J. Differential Geom.*, 73(3):363–394, 2006.

[24] K. Crane, U. Pinkall, and P. Schröder. Robust fairing via conformal curvature flow. *ACM Trans. Graph.*, 32, 2013.

[25] B. Dacorogna. *Direct methods in the calculus of variations*. Springer-Verlag, New York, 1989.

[26] M. Desbrun, A. Hirani, M. Leok, and J. E. Marsden. Discrete exterior calculus. 2005. `arXiv:math.DG/0508341`.

[27] M. Desbrun, M. Meyer, P. Schröder, and A. Barr. Implicit fairing of irregular meshes using diffusion and curvature flow. In *Computer Graphics (SIGGRAPH '99 Proceedings)*, pages 317–324, 1999.

[28] G. Dziuk. *Finite elements for the Beltrami operator on arbitrary surfaces*, volume 1357 of *Lecture Notes in Math*, chapter Partial differential equations and calculus of variations, pages 142–155. Springer, 1988.

[29] G. Dziuk. An algorithm for evolutionary surfaces. *Numer. Math.*, 58:603–611, 1991.

[30] G. Dziuk and C. M. Elliott. Surface finite elements for parabolic equations. *Journal of Computational Mathematics*, 25:385–407, 2007.

[31] G. Dziuk and C. M. Elliott. L^2-estimates for the evolving surface finite element method. *Math. Comp.*, 82(281):1–24, 2013.

[32] H. Edelsbrunner. *Geometry and topology for mesh generation*. Cambridge University Press, 2001.

[33] P. Fletcher, C. Lu, S. Pizer, and S. Joshi. Principal geodesic analysis for the study of nonlinear statistics of shape. *Medical Imaging, IEEE Transactions on*, 23(8):995–1005, 2004.

[34] M. S. Floater and K. Hormann. Surface Parameterization: A Tutorial and Survey. In N. A. Dodgson, M. S. Floater, and M. A. Sabin, editors, *Advances in Multiresolution for Geometric Modeling*, Mathematics and Visualization, pages 157–186. Springer, 2005.

[35] G. Friesecke, R. James, and S. Müller. A theorem on geometric rigidity and the derivation of nonlinear plate theory from three-dimensional elasticity. *Comm.Pure Appl.Math.*, 55(11):1461–1506, 2002.

[36] M. Fuchs and O. Scherzer. Regularized reconstruction of shapes with statistical a priori knowledge. *International Journal of Computer Vision*, 79(2):119–135, 2008.

[37] E. Grinspun, Y. Gingold, J. Reisman, and D. Zorin. Computing discrete shape operators on general meshes. *Eurographics (Computer Graphics Forum)*, 25, 2006.

[38] E. Grinspun, A. N. Hirani, M. Desbrun, and P. Schröder. Discrete shells. In *Siggraph/Eurographics Sympos. Comput. Anim.*, pages 62–67, 2003.

[39] E. Grinspun, P. Krysl, and P. Schröder. CHARMS: A Simple Franmework for Adaptive Simulation. In *Computer Graphics (SIGGRAPH '02 Proceedings)*, 2002.

[40] D. Harmon, E. Vouga, B. Smith, R. Tamstorf, and E. Grinspun. Asynchronous contact mechanics. *Communications of the ACM*, 55(4):102–109, 2012.

[41] M. Hauth. *Visual Simulation of Deformable Models*. PhD thesis, University of Tübingen, 2004.

[42] K. Hildebrandt and K. Polthier. Anisotropic filtering of non-linear surface features. *Eurographics*, 23(3), 2004.

[43] K. Hildebrandt and K. Polthier. Generalized shape operators on polyhedral surfaces. *Computer Aided Geometric Design*, 28(5):321–343, 2011.

[44] K. Hildebrandt and K. Polthier. On approximation of the Laplace–Beltrami operator and the Willmore energy of surfaces. *Computer Graphics Forum*, 30(5):1513–1520, 2011.

[45] K. Hildebrandt, K. Polthier, and M. Wardetzky. On the convergence of metric and geometric properties of polyhedral surfaces. *Geometricae Dedicata*, 123:89–112, 2006.

[46] K. Hildebrandt, C. Schulz, C. V. Tycowicz, and K. Polthier. Interactive surface modeling using modal analysis. *ACM Transactions on Graphics*, 30(5):119:1–119:11, 2011.

[47] B. Houston, M. Nielsen, C. Batty, O. Nilsson, and K. Museth. Hierarchical RLE level set: A compact and versatile deformable surface representation. *ACM Transactions on Graphics*, 25(1):1–24, 2006.

[48] T. J. R. Hughes. *Finite Element Method - Linear Static and Dynamic Finite Element Analysis*. Prentice-Hall, Englewood Cliffs, 1987.

[49] G. Huisken. The volume preserving mean curvature flow. *Journal für die reine und angewandte Mathematik*, 382:35–48, 1987.

[50] J. Iglesias, B. Berkels, M. Rumpf, and O. Scherzer. A shell model for level set matching. in preparation.

[51] J. Iglesias, B. Berkels, M. Rumpf, and O. Scherzer. A thin shell approach to the registration of implicit surfaces. In M. Bronstein, J. Favre, and K. Hormann, editors, *Vision, Modeling and Visualization*, pages 89–96. Eurographics Association, 2013.

[52] Z. Karni and C. Gotsman. Spectral compression of mesh geometry. In *Proceedings of ACM SIGGRAPH*, pages 279–286. ACM Press, 2000.

[53] B. Kawohl and N. Kutev. Maximum and comparison principle for one-dimensional anisotropic diffusion. *Math. Ann.*, 311(1):107–123, 1998.

[54] M. Kazhdan, M. Bolitho, and H. Hoppe. Poisson surface reconstruction. In *Proceedings of the Symposium of Geometry Processing*, pages 61–70, 2006.

[55] M. Kilian, N. J. Mitra, and H. Pottmann. Geometric modeling in shape space. In *ACM Transactions on Graphics*, volume 26, pages 1–8, 2007.

[56] E. Klassen, A. Srivastava, W. Mio, and S. H. Joshi. Analysis of planar shapes using geodesic paths on shape spaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(3):372–383, 2004.

[57] A. Lew, M. Marsden, M. Oritz, and M. West. Variational time integrators. *Int. J. Numer. Meth. Engng*, 60:153–212, 2004.

[58] Y. Lipman. Bounded distortion mapping spaces for triangular meshes. *ACM Transactions on Graphics*, 31(4):108, 2012.

[59] Y. Lipman, O. Sorkine, D. Levin, D. Cohen-Or, C. Rössl, and H.-P. Seidel. Differential coordinates for interactive shape editing. *Proceedings of SMI*, pages 181–190, 2004.

[60] N. Litke, M. Droske, M. Rumpf, and P. Schröder. An image processing approach to surface matching. In M. Desbrun and H. Pottmann, editors, *Third Eurographics Symposium on Geometry Processing*, pages 207–216, 2005.

[61] C. Loop. Smooth spline surfaces over irregular meshes. In *SIGGRAPH '94: Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 303–310. ACM Press, 1994.

[62] R. MacNeal. The solution of partial differential equations by means of electrical networks, 1949. PhD thesis, California Institute of Technology.

[63] C. Mandal, H. Qin, and B. Vemuri. A novel fem-based dynamic framework for subdivision surfaces. *Computer-Aided Design*, 32(8-9):479–497, 2000.

[64] P. W. Michor and D. Mumford. Riemannian geometries on spaces of plane curves. *J. Eur. Math. Soc.*, 8:1–48, 2006.

[65] M. I. Miller, A. Trouvé, and L. Younes. The metric spaces, Euler equations, and normal geodesic image motions of computational anatomy. In *Proceedings of the 2003 International Conference on Image Processing*, volume 2, pages 635–638. IEEE, 2003.

[66] J.-M. Morvan. *Generalized Curvatures*. Geometry and Computing. Springer, 2008.

[67] P. Mullen, Y. Tong, P. Alliez, and M. Desbrun. Spectral conformal parameterization. *Computer Graphics Forum (Proc. of SGP)*, 27(5):1487–1494, 2008.

[68] A. Nealen. An as-short-as-possible introduction to the moving least squares method for scattered data approximation and interpolation, 2004. Technical Report.

[69] O. Nemitz, M. B. Nielsen, M. Rumpf, and R. Whitaker. Finite element methods on very large, dynamic tubular grid encoded implicit surfaces. *SIAM Journal on Scientific Computing*, 31(3):2258–2281, 2009.

[70] S. Osher and R. P. Fedkiw. Level set methods: An overview and some recent results. *Journal of Computational Physics*, 169(2):463 – 502, 2001.

[71] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7):629–639, July 1990.

[72] U. Pinkall and K. Polthier. Computing discrete minimal surfaces and their conjugates. *Experim. Math.*, 2:15–36, 1993.

[73] T. Preußer and M. Rumpf. A level set method for anisotropic geometric diffusion in 3D image processing. *SIAM Journal on Applied Mathematics*, 62(5):1772–1793, 2002.

[74] L. Rayleigh. *Theory of Sound: V. 2, Band 2*. Courier Dover Publications, 1900.

[75] M. Reuter. Hierarchical shape segmentation and registration via topological features of laplace-beltrami eigenfunctions. *International Journal of Computer Vision*, 89(2):287–308, 2010.

[76] M. Reuter, F.-E. Wolter, and N. Peinecke. Laplace-beltrami spectra as "shape-dna" of surfaces and solids. *Computer-Aided Design*, 38(4):342–366, 2006.

[77] M. Rumpf and B. Wirth. Variational time discretization of geodesic calculus, 2012. submitted.

[78] M. Rumpf and B. Wirth. Discrete geodesic calculus in the space of viscous fluidic objects. *SIAM Journal on Imaging Sciences*, 2014. in press.

[79] J. A. Sethian. *Level Set Methods: Evolving Interfaces in Geometry, Fluid Mechanics, Computer Vision and Materials Sciences*. Cambridge Univ. Press, 1996.

[80] J. A. Sethian. *Level set methods and fast marching methods*, volume 3 of *Cambridge Monographs on Applied and Computational Mathematics*. Cambridge University Press, Cambridge, second edition, 1999. Evolving interfaces in computational geometry, fluid mechanics, computer vision, and materials science.

[81] O. Sorkine and M. Alexa. As-rigid-as-possible surface modeling. In *Proceedings of EUROGRAPHICS/ACM SIGGRAPH Symposium on Geometry Processing*, pages 109–116, 2007.

[82] O. Sorkine, D. Cohen-Or, Y. Lipman, M. Alexa, C. Rössl, and H.-P. Seidel. Laplacian surface editing. In *Proceedings of the EUROGRAPHICS/ACM SIGGRAPH Symposium on Geometry Processing*, pages 179–188. ACM Press, 2004.

[83] A. Srivastava, E. Klassen, S. H. Joshi, and I. H. Jermyn. Shape analysis of elastic curves in euclidean spaces. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 33(7):1415–1428, july 2011.

[84] J. Sun, M. Ovsjanikov, and L. Guibas. A concise and provably informative multi-scale signature based on heat diffusion. In *Eurographics Symposium on Geometry Processing (SGP)*, 2009.

[85] G. Sundaramoorthi, A. Mennucci, S. Soatto, and A. Yezzi. A new geometric metric in the space of curves, and applications to tracking deforming objects by prediction and filtering. *SIAM Journal on Imaging Sciences*, 4(1):109–145, 2011.

[86] G. Sundaramoorthi, A. Yezzi, and A. Mennucci. Sobolev active contours. *International Journal of Computer Vision.*, 73(3):345–366, 2007.

[87] D. Terzopoulos, J. Platt, A. Barr, and K. Fleischer. Elastically deformable models. In *Proceedings of SIGGRAPH*, pages 205–214, 1987.

[88] G. Turk and J. O'Brien. Modelling with implicit surfaces that interpolate. *ACM Transactions on Graphics*, 21:855–873, 2002.

[89] C. von Tycowicz, C. Schulz, H.-P. Seidel, and K. Hildebrandt. An efficient construction of reduced deformable objects. *ACM Trans. Graph.*, 32(6):213:1–213:10, Nov. 2013.

[90] M. Wardetzky, M. Bergou, D. Harmon, D. Zorin, and E. Grinspun. Discrete Quadratic Curvature Energies. *Computer Aided Geometric Design*, 24:499–518, 2007.

[91] M. Wardetzky, S. Mathur, F. Kälberer, and E. Grinspun. Discrete Laplace operators: No free lunch. In *Siggraph/Eurographics Sympos. Geom. Processing*, pages 33–37, 2007.

[92] J. Weickert. *Anisotropic diffusion in image processing*. Teubner, 1998.

[93] O. C. Zienkiewicz and R. L. Taylor. *The finite element method: The basis*, volume 1. Butterworth and Heinemann, 5th edition, 2000.