Geometry Processing from an Elastic Perspective

Martin Rumpf[1,*] and **Max Wardetzky**[2,**]

[1] Institut for Numerical Simulation, University of Bonn, Endenicher Allee 60, 53115 Bonn, Germany

[2] Institute of Numerical and Applied Mathematics, University of Göttingen, Lotzestr. 16-18, 37083 Göttingen, Germany

Received XXXX, revised XXXX, accepted XXXX
Published online XXXX

Key words geometric partial differential equations, geometric functionals, thin shells, shape spaces
MSC (2010) 68U10, 53C22, 74B20, 49M20

Triggered by the development of new hardware, such as laser range scanners for high resolution acquisition of complex geometric objects, new graphics processors for realtime rendering and animation of extremely detailed 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 of both fundamental research and impressive applications. Concepts from image processing have been picked up and carried over to curved surfaces, physically based modeling plays a central role, and aspects of computer aided geometry design have been incorporated. This paper aims at highlighting some of these developments, with a particular focus on methods related to the mechanics of thin elastic surfaces. We provide an overview of different geometric representations ranging from polyhedral surfaces over level sets to subdivision surfaces. Furthermore, with an eye on differential-geometric concepts underlying continuum mechanics, we discuss fundamental computational tasks, such as surface flows and fairing, surface deformation and matching, physical simulations, as well as spectral and modal methods in geometry processing. Finally, beyond focusing on single shapes, we describe how spaces of shapes can be investigated using concepts from Riemannian geometry.

1 Introduction

Over the last decades, surface models used in computer graphics have become successively more complex. New hardware, such as laser range scanners for the acquisition of 3D geometries, and new software, such as 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 application, different geometric representations are appropriate to process, model, or animate the underlying surface geometries. On this background, the research field of geometry processing has undergone a rapid development. In particular, physically based

* E-mail: martin.rumpf@uni-bonn.de,
** E-mail: wardetzky@math.uni-goettingen.de

modeling plays a crucial role, and concepts from continuum mechanics have been picked up and adapted to the needs in computer graphics.

One overarching goal of geometry processing is the combination of valid physical models with efficient, near realtime simulations. Recently, it has become an established trend to *not* sacrifice physical accuracy for computational efficiency. One important instance of this trend is to work with reliable nonlinear models and to use geometric and 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 another. E.g., when working with triangle meshes (or other discrete geometric representations of smooth surfaces) a central goal is to provide structure preserving, consistent, and convergent notions 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 opts for an implicit representation of surfaces via level set methods, the computational workload significantly raises from dimension two to three.

This paper highlights some of the ongoing developments in geometry processing, with a particular focus on methods related to and motivated by the mechanics of thin shells. As a disclaimer, we 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. 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 existing in the literature.

In our exposition, we review different representations of surfaces, ranging from parameterized and triangulated surfaces over subdivision surfaces to level sets and point clouds in Section 2. 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 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 parameterized and implicit surfaces and discuss discrete analogous of these energies using discrete exterior calculus. Applications to surface matching and physical simulations are outlined. The spectral analysis of the Hessian of elastic energies and the computation of eigenfrequencies and vibration modes 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 (i.e., a space in which each point is a surface in $\mathbb{R}^3$), and we sketch the setup of a time discrete Riemannian calculus in this space. Physically, the underlying Riemannian metric that we consider represents the rate of physical dissipation accumulated along paths in shell space.

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 applications, different tools are implemented on different types

of geometry representations. To emphasize the similarities and relations between these approaches, we discuss in this section some of the most basic differential operators and energy functionals relevant to geometry processing, namely the Dirichlet energy, the Laplacian, and the Willmore functional. Furthermore, we briefly sketch how these operators and functionals can be spatially discretized.

The flat case. We briefly recall the Euclidean case. On a two dimensional domain $\Omega \subset \mathbb{R}^2$ the Dirichlet energy of a function $u : \Omega \to \mathbb{R}$ is defined as

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

where ∇ denotes the gradient operator. The weak form of the (negative) Laplace operator $-\triangle = -\mathrm{div}\nabla$ applied to a function u is then given as the first variation of Dirichlet energy.[1] Indeed, using integration by parts one easily verifies that

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

for a sufficiently regular function u and for any smooth test function ϕ that is compactly supported away from the boundary of Ω.

A simple model for a thin plate energy on the planar domain Ω and for a vertical displacement u is given by $\mathcal{W}_{\mathrm{plate}}[u] = \frac{1}{2}\int_\Omega |\triangle u|^2 \, \mathrm{d}x$. The first variation of this energy is given using the bi-Laplacian as $\int_\Omega \triangle^2 u \phi \, \mathrm{d}x = \int_\Omega \triangle u \cdot \triangle \phi \, \mathrm{d}x$, 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 $w = -\triangle u$ 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) triangle 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 $T \in \mathcal{T}_h$. 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 $\mathcal{W}_{\mathrm{Dirichlet}}$ on discrete functions U gives rise to a bilinear 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 $\mathbf{M}\bar{U} \cdot \bar{V} := \int_\Omega \mathcal{I}_h(U \cdot V) \, \mathrm{d}x$ (where $\mathcal{I}_h$ denotes the piecewise affine Lagrangian interpolation of the vertex values of a function on $\mathcal{T}_h$), one denotes by $-\triangle_h \bar{U} = \mathbf{M}^{-1}\mathbf{L}\bar{U}$ the discrete negative Laplacian on nodal vectors—in analogy to the above relation for the continuous Laplacian (1). 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 $\mathcal{W}^h_{\mathrm{plate}}[U] = \frac{1}{2}\int_\Omega (\mathcal{I}_h(\triangle_h(\bar{U})))^2 \, \mathrm{d}x$.

In what follows, we adopt this calculus for a geometric setting.

[1] Recall that the first variation of a nonlinear functional of u yields a first order approximation, i.e., a linear functional with parameter u.

Smooth embedded surfaces. Let $\mathcal{S} \subset \mathbb{R}^3$ be a smoothly embedded and oriented surface. Let n denote the unit normal. Then the *first and second fundamental forms* on $\mathcal{S}$ are bilinear 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 curvature* K and the *mean curvature* h are defined as the determinant and trace of $\mathbf{S}$, respectively.[2]

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 denotes the area element of $\mathcal{S}$. With these definitions, one can define the Dirichlet energy of functions $\mathcal{W}_{\mathrm{Dirichlet}}[f] = \frac{1}{2}\int_{\mathcal{S}} |\nabla_{\mathcal{S}} f|^2 \,\mathrm{d}a$, 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.,

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

for smooth functions $\phi : \mathcal{S} \to \mathbb{R}$ that are compactly supported away from the boundary of $\mathcal{S}$. Instead of scalar functions, we can consider vector valued functions and apply differential operators to all components of the function. With a slight abuse of notation, we also use x to denote the vector valued mapping $x : \mathcal{S} \to \mathbb{R}^3$, which maps every point x onto itself. The fundamental geometric insight is that $\triangle_{\mathcal{S}} x = -hn$, where hn is the mean curvature vector. The 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 inner product. Hence, the associated gradient flow of the area functional is the geometric heat equation $\partial_t x - \triangle_{\mathcal{S}} x = 0$. The nonlinear counterpart of the thin plate energy is the Willmore energy $\mathcal{W}_{\mathrm{Willmore}}[x] = \frac{1}{2}\int_{\mathcal{S}} h^2 \,\mathrm{d}a$.[3] 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. Using this representation, 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 forms 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 local parameterizations.

[2] Notice that by defining mean curvature as $h = \mathrm{tr}\mathbf{S}$, we deviate from other authors that use $h = \frac{1}{2}\mathrm{tr}\mathbf{S}$.

[3] Notice that as long as the surface boundary is held fixed, our formulation of Willmore energy has the same critical points as the Möbius-invariant formulation $\int_{\mathcal{S}} (\kappa_1 - \kappa_2)^2 \,\mathrm{d}a$.

Local parameterizations. Given a surface $\mathcal{S} \subset \mathbb{R}^3$, one *cannot* in general obtain a *global parameterization*, i.e., a diffeomorphism from a region in the plane to $\mathcal{S}$, which 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 parameterizations, however, requires to account for coordinate changes, which is a tedious task in practice. Therefore, it is often desirable to work with global representations of surfaces, such as triangle meshes or level set representations, which do not require any particular parameterization. Moreover, the *intrinsic*, i.e., parameterization independent, formulation of differential operators is often more natural and convenient. Nonetheless, various applications, such as texture mapping, do require local parameterizations, and higher order discretization methods such as subdivision finite elements use local parameterizations over control meshes as well. There is an extensive literature (see, e.g., [1] and references therein) on the computation of local parameterizations, dealing in particular with the issue of minimizing *distortion of lengths, angles, and area*. Since distortion is unavoidable in general when mapping a nonplanar to a planar domain, this poses the question if one can at least bound the amount of distortion introduced. For the case of triangle meshes and piecewise linear continuous bijections to a planar region, this problem has recently been solved by Lipman for the case of aspect ratio distortion [2], see Figure 1. A detailed discussion of local parameterizations is beyond the scope of this paper.

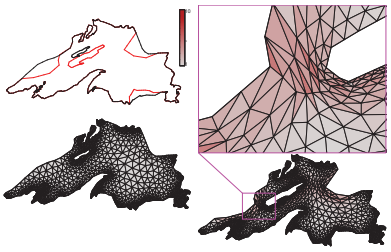

Fig. 1 Bounded distortion mapping of Lake Superior. Original mesh (left) is deformed by moving island and boundaries. Bounded distortion maps (right) avoid triangle flips and high distortion. (Image courtesy of Y. Lipman.)

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 $n(x) = \partial_{\xi_1} x \times \partial_{\xi_2} x / \|\partial_{\xi_1} x \times \partial_{\xi_2} x\|$ 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 is $g^{-1} = (g^{ij})_{i,j=1,2}$. With the metric at hand, the integral of a function f on $x(\Omega)$ is given by $\int_{x(\Omega)} f \,\mathrm{d}a = \int_\Omega f \circ x \sqrt{g}\,\mathrm{d}\xi$. 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 $\mathbf{S}^\Omega = g^{-1} Dx^T D(n \circ x)$ as a 2×2 matrix defined on the chart. The above mentioned differential operators can also 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 $(\mathrm{div}_\mathcal{S} \mathbf{v}) \circ x = \frac{1}{\sqrt{g}} \mathrm{div}_\Omega (v\sqrt{g})$. Hence, the Laplace Beltrami operator is given by $\triangle_\mathcal{S} u = \mathrm{div}_\mathcal{S} \nabla_\mathcal{S} u = \frac{1}{\sqrt{g}} \mathrm{div}_\Omega \left(g^{-1} \nabla_\Omega (u \circ x) \sqrt{g} \right)$.

Triangulated surfaces. Following the discretization procedure in the flat Euclidean case, we consider a triangulated surface $\mathcal{S}_h$ consisting of flat triangles. We denote by $\mathcal{V}_h$ the

space of piecewise affine, continuous functions on $\mathcal{S}_h$. As before, the variation of the discrete Dirichlet energy $\mathcal{W}^h_{\text{Dirichlet}} = \frac{1}{2}\int_{\mathcal{S}_h} |\nabla_{\mathcal{S}_h} U|^2 \, \mathrm{d}a$ for functions $U \in \mathcal{V}_h$ yields a bilinear 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 $\bar{V}$ with 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 and mass matrix can be computed as follows. 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 $(\mathbf{L}_h^{\text{local}})_{ij} = \frac{1}{4} e_i \cdot e_j |T|^{-1}$ and the local (lumped) mass matrix $(\mathbf{M}_h^{\text{local}})_{ij} = \frac{1}{3}|T|\delta_{ij}$, which have to be assembled to their global counterparts in the usual way. Notice that $(\mathbf{L}_h^{\text{local}})_{ij} = -\frac{1}{2}\cot(\alpha_{ij}))$, 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* [3, 4]. Analogously, we define the lumped mass matrix $\mathbf{M}\bar{U} \cdot \bar{V} := \int_{\mathcal{S}_h} \mathcal{I}_h(U \cdot V) \, \mathrm{d}a$ 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 $\bar{U}$, see [5]. *Convergence* of this approach for polyhedral surfaces that approximate (but do not necessarily interpolate) a smooth limit surface has been established in [6].

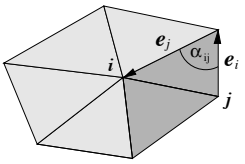

Discrete curvatures and first and second fundamental forms. One challenge for polyhedral surfaces is to provide consistent notions of curvatures and second fundamental forms, i.e., notions that converge (in an appropriate topology or in a measure theoretic sense) to their smooth counterparts, given a smooth limit surface $\mathcal{S} \subset \mathbb{R}^3$.

One computationally popular model for discretizing the second fundamental form is Grinspun's et al. *discrete shells* model [7], which offers a discrete version of *mean curvature* associated with edges. In this view, mean curvature is measured as the (signed) dihedral angle of a an edge, multiplied by the length of the edge. This view corresponds to an elegant mathematical model—the mean curvature *measure* arising form the theory of *normal cycles* [8]. Normal cycles generalize the notion of unit normal bundles of oriented, smooth submanifolds to more general cases, such as (boundaries of) convex bodies and closed polyhedral surfaces. E.g., for a convex set $S \subset \mathbb{R}^3$, the notion of a unit normal at a point $p \in S$ is replaced by the normal cone $C_p = \{n \in \mathbb{R}^3 \,|\, \|n\| = 1 \,,\, n \cdot (q - p) \leq 0 \, \forall q \in S\}$. The normal cycle $N(S)$ of a convex body S is then associated with the set $\{(p, n) \,|\, p \in S \,,\, n \in C_p\} \subset \mathbb{R}^3 \times \mathbb{S}^2$. In particular, this construction applies to tetrahedra in $\mathbb{R}^3$. To define the normal cycle of the geometric realization of a 3-dimensional simplicial complex in $\mathbb{R}^3$, one starts from the definition of normal cycles for 3-simplices (tetrahedra) and successively builds up the normal cycle of the entire complex by requiring the inclusion-exclusion property $N(A) + N(B) = N(A \cup B) + N(A \cap B)$.

Much like for the case of embedded smooth surfaces, one can define an area measure as well as measures for mean and Gauß curvatures for normal cycles, giving rise to discrete curvatures of polyhedral surfaces. In particular, discrete Gauß curvature, associated with mesh vertices, coincides with the widely used *angle deficit*, i.e., $K_p = 2\pi - \sum_i \alpha_i$, where α_i are the interior angles at a vertex p of the triangles meeting at p. Moreover, Cohen–Steiner and Morvan used the theory of normal cycles to provide a notion of discrete shape operators for polyhedral surfaces [9, 10]. This notion turns out to converge in the sense of measures. A similar notion that also converges in the sense of measures, leading to a consistent notion

of the Willmore functional on polyhedral surfaces, has been proposed by Hildebrandt and Polthier [11]. 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. Indeed, convergence can be shown by letting the averaging region shrink at a *much* lower rate than the refinement of triangles increases within that region. The averaging effect is much more pronounced for unstructured meshes (e.g., Delaunay triangulation) than it is for structured meshes, see [12].

A slight modification of these approaches leads to the following formulation of discrete second fundamental forms. We elaborate on this construction, since it leads (i) to a formulation of second fundamental forms that is *constant* per triangle and (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 $\tilde{T}$, and associate n_e with the *midpoint* of e. For boundary edges, consider the respective triangle normal. With normals associated to edge midpoints, one obtains a vector-valued discrete 1-form $\mathrm{d}n$ (i.e., a linear map from oriented line segments to $\mathbb{R}^3$). Indeed, 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 bilinear form is *uniquely determined* by its action on three different vectors, leads to a discrete second fundamental form that is constant per triangle:

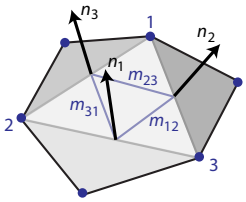

$$\mathrm{II}_T = \frac{1}{8|T|^2} \sum_{i=1}^{3} (\mathrm{II}_{T,j} + \mathrm{II}_{T,k} - \mathrm{II}_{T,i})\, t_i \otimes t_i \,, \tag{4}$$

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 $\mathrm{I}_{T,k} := \mathrm{I}_T(e_k, e_k) := |e_k|^2$ 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 of the first fundamental form is known as the *constant strain triangle*.

Subdivision Finite Elements. As long as smooth surfaces are approximated by triangular meshes, 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

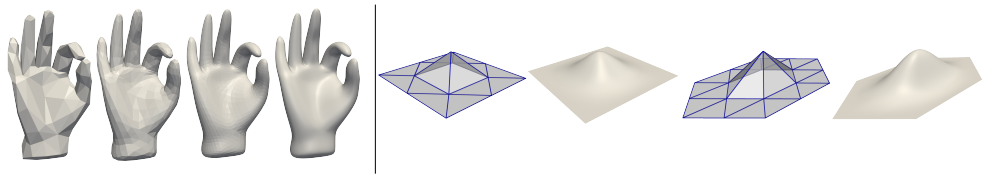

Fig. 2 The first two subdivision steps starting from the control mesh and the resulting limit surface are visualized (left). Furthermore, subdivision bases functions are rendered for a node with valence 4 (left) and 6 (right), respectively. (Image courtesy of R. Perl.)

triangular mesh could be considered as a control mesh of a much smoother surface representation. Such a representation can be generated 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, followed by replacing every vertex of the refined mesh by a weighted sum of vertices in the support mask of a given one. In their pioneering work, Catmull and Clark [13] 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 [14] (see Fig. 2). If restricted to control patches with valence six, i.e., a situation where every vertex is shared by six adjacent triangles, Loop's subdivision scheme is equivalent to the construction of bi-variate box-splines, and thus the limit surface is C^2; 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 $C^1 \cap H^2$, and thus the mean curvature is still in L^2. Mandal, Qin, and Vemuri [15] used subdivision surfaces for 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 parameterization over flattened control meshes. They can be applied to general scalar values or vectors assigned to the control vertices in order 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 bi-Laplacian on the limit surface $\mathcal{S}_h^\infty$ of a control mesh $\mathcal{S}_h$ is given by

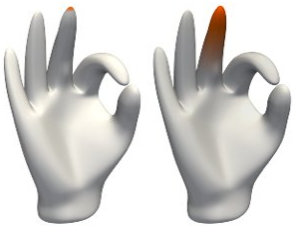

Fig. 3 The equation $\triangle_S^2 u + u = f$ is solved on a surface using subdivision finite elements based on Loop subdivision with color coded f (left) and u (right). (Image courtesy of R. Perl.)

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

for a subdivision finite element function u and all subdivision finite element test functions ϕ (see Eq. (2)). Figure 3 shows the result for solving the equation $\triangle_S^2 u + u = f$ on a surface using subdivision finite elements. Cirak, Ortiz, and Schröder [16] investigated a thin-shell energy using subdivision finite elements to describe the surface geometry and to compute smooth displacement fields in a conforming finite element Galerkin approach. They applied

this approach to the Kirchhoff–Love theory of thin shells using Loop's subdivision scheme. In this approach, for the assembly of finite element matrices on valence six control patches, the explicit spline representation can be retrieved and the computations can be performed on the control mesh as a parameter domain of the actual subdivision limit surface. In the case of general valences, one uses an appropriate quadrature formula with only interior quadrature points and applies the subdivision scheme recursively until all quadrature points lie in a valence six patch and thus the explicit representation again applies (see Fig. 2). In [17] Grinspun, Krysl, and Schröder proposed an adaptive subdivision finite element scheme in order to accelerate computations.

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 one obtains $h = \operatorname{tr}(DnP) = \operatorname{div} n$. 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, one does usually not perform integration over a single surface; instead, one integrates over a bundle of surfaces $[c^- < w < c^+] = \bigcup_{c\in[c^-,c^+]} \mathcal{S}^c$, in which case the coarea formula implies that

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

with $|\nabla w|$ representing the density of level sets. As a consequence, 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 as

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

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

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

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

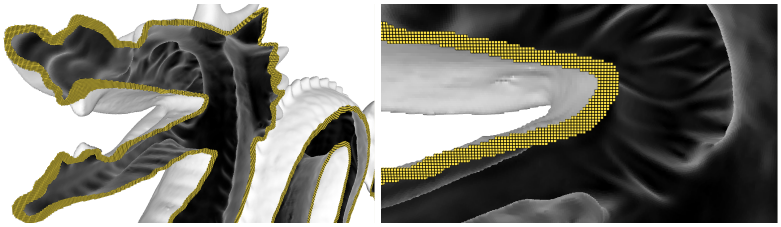

Fig. 4 Level set representation of a complex dragon model on a thin narrow band using a cutting plane intersecting the dragon surface. The underlying spatial resolution is $1986 \times 1323 \times 1104$.

$\partial_t w + |\nabla w| \mathrm{div}(|\nabla w|^{-1} \nabla w) = 0$. Finally, the Willmore energy integrated over a bundle of surfaces $\{\mathcal{S}_c \,|\, c \in [c^-, c^+]\}$ is given by

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

We emphasize that in the level set formulation, PDE problems on different surfaces $\mathcal{S}_c$ and $\mathcal{S}_{\bar{c}}$ 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 $\Omega \subset \mathbb{R}^3$, 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, as before, $\mathcal{I}_h$ 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 [18]. For the general concepts of the level set method we refer to the textbooks by Sethian [19] or Osher and Fedkiw [20]. 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 neighborhood 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 [21]. Fig. 4 depicts the narrow band geometry for a highly detailed 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, [22–24]. Other approaches base reconstruction on implicit surfaces, e.g., by using radial basis functions (RBFs) or moving least squares [25, 26]. 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 [27] (see Figure 5). In summary, these methods transform point clouds into surface representations

Fig. 5 *Poisson Surface Reconstruction* proposed in [27] for high quality surface reconstruction from oriented points clouds. (Image courtesy of M. Kazhdan.)

discussed above, whereafter one works with the tools available for these representations for solving PDEs or formulating geometric energy functionals. Another approach is to avoid a transformation step altogether, such as, e.g., in Belkin's et al. work [28], where the authors use short-time properties of the heat kernel in $\mathbb{R}^3$ for constructing Laplacians on point clouds.

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 an abstract manifold $\mathcal{M}$. 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 form, 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 properties of these operators, we content ourselves with pointing out that exterior differentiation is a generalization 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 $\mathcal{M}$. Equipped with a Riemannian metric, one defines 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 Hodge 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) [29]. 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 then canonically taken by the simplicial coboundary operator $\delta : C^k \to C^{k+1}$. For constructing a discrete Hodge star operator, Hirani et al. consider the *circumcentric* dual complex K^* of a simplicial complex K. This dual construction gives rise to an operator $* : C_k \to C_{n-k}$ taking primal chains in K to dual chains in K^* and vice versa. The discrete Hodge star $\star : C^k \to C^{n-k}$ is then defined as a (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 $\triangle := \delta^*\delta + \delta\delta^*$. For the Laplacian acting on 0-cochains (i.e., discrete functions defined at vertices), one exactly recovers the cotangent operator from the FE setting described above. An alternative approach for defining δ^* (and therefore $\triangle$) is, as the the smooth case, to consider L^2-inner products $(\cdot,\cdot)_{L^2}$ on k-cochains and to define 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 the inner product approach additionally leads to a construction of Laplacians on discrete surfaces with general polygonal (not necessarily triangular) faces [30]. The diagonal Hodge star approach of DEC, in turn, can by utilized for mesh optimization and for improving numerical accuracy by working with Hodge stars resulting from dual meshes *different* from circumcentric duals—see the work of de Goes, Mullen, Desbrun, et al. [31, 32].

In a similar spirit to DEC, Arnold, Falk, and Winther have developed *Finite Element Exterior Calculus* (FEEC) [33]. One of the main principles 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 $(\omega,\tau)_{L^2} = (\sigma, d\tau)_{L^2}$, leading to a *mixed* formulation for the Hodge Laplacian. For piecewise linear functions, DEC and FEEC share various similarities.

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 smooth surfaces, such as minimal surfaces, surfaces of constant mean curvature, or surfaces of constant negative Gauß curvature, is governed by nonlinear PDEs, such as the KdV equation, that give rise to completely integrable systems, and (one branch of) the field of *Discrete Differential Geometry* (DDG) searches for discretizations that preserve the underlying (integrable) structure of the smooth case. For readers interested in this aspect, we refer to [34–36].

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 in [37] that there exists no discrete Laplacian that possesses *all* properties of the smooth one. Indeed, the main hurdle 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 coarser for $t \to \infty$, thereby successively smoothing out the initial data.

Isotropic diffusion / mean curvature flow. It is near at hand to ask for analogous strategies for the fairing of perturbed and noisy surface geometries $\mathcal{S}_0$, and this naturally leads to mean curvature motion $\partial_t x - \triangle_{\mathcal{S}} x = 0$ using the fact that $hn = -\triangle_{\mathcal{S}} x$. There exist numerous implementations of this flow, going back at least to Brakke's *surface evolver* [38]. Discretizing this evolution problem explicitly in time, 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 $\mathcal{S}^n$. In order to speed up this scheme and to improve its robustness, Desbrun et al. [39] considered the implicit scheme $x^{n+1} = x^n - \tau \triangle_{\mathcal{S}^n} x^{n+1}$ (see Fig. 6), similar to the one proposed by Dziuk [40], where the Laplace–Beltrami operator is still evaluated at the old time step. The mean curvature motion model is known as the L^2 gradient flow of the surface area $\mathcal{A}(\mathcal{S})$, and $\frac{d}{dt}\mathcal{A}(\mathcal{S}(t)) = -\int_{\mathcal{S}(t)} h^2 \, da$ (see [41]), 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 acquisition, but it also smoothes out geometric features, such as edges and corners. Hence, a model is required that improves a simple high pass filtering. Here image processing methodology can be used.

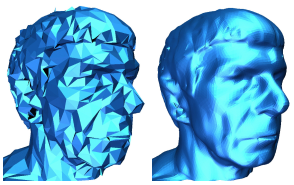

Fig. 6 The smoothing of an initial surface (left) by time discrete mean curvature motion (right) using the implicit fairing approach by Desbrun et al. [39]. (Image courtesy of M. Desbrun.)

Anisotropic diffusion. For the case of images, Perona and Malik [42] proposed a nonlinear diffusion method based on the PDE $\partial_t u - \operatorname{div}(a(|\nabla u|)\nabla u) = 0$ with a diffusion coefficient $a(s) = \left(1 + \frac{s^2}{\lambda^2}\right)^{-1}$ that 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 (see [43]). To overcome the ill-posedness of the original Perona–Malik problem, Catté et al. [44] proposed a regularization method, where the diffusion coefficient $a(\cdot)$ is evaluated on a prefiltered image intensity $u_\sigma = G_\sigma * u$. Weickert [45] 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.

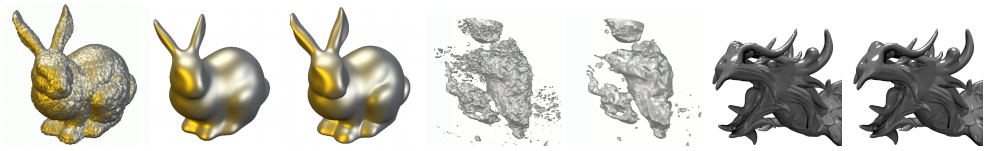

Fig. 7 Mean curvature motion fairing is compared with anisotropic geometric diffusion for a noisy triangular surface (left), to a level set of a 3D echocardiographical image of a ventricle of the human heart (middle) and to a a noisy level set model of a dragon using a narrow band approach for high resolution surfaces (right).

Anisotropic diffusion of 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 $\mathcal{S}_\sigma$, that is $A^\sigma v_\sigma^i \cdot v_\sigma^j = \delta_{ij} a(\kappa_\sigma^i)$ with κ_σ^i denoting the corresponding principal curvatures. Following the general discretization procedure outlined in Section 2, the method can be discretized on triangular meshes using affine finite elements. A similar approach has been investigated by Hildebrandt and Polthier [46]. 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 [47]. Fig. 7 shows results of the anisotropic diffusion method.

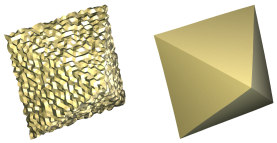

Fig. 8 Result of the discrete geometry approach to anisotropic diffusion (right) for a noisy octahedron (left). (Image courtesy of K. Hildebrandt.)

Anisotropic diffusion of implicit surfaces / level sets. The above approach can be adapted to the processing of level set surfaces. The advantage of level set formulations is their ability to deal with *pinching off* and topology changes due to singularities that the flow can develop in finite time. Denote by $w_0 : \Omega \to \mathbb{R}$ 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 that solves the anisotropic diffusion problem

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

with natural (no flux) boundary conditions on $\partial\Omega$ and initial data $w(0) = w_0$. In this case $A^\sigma = a(\mathbf{S}_\sigma^{\text{ext}}) - n_\sigma \otimes n_\sigma$ with $\mathbf{S}_\sigma^{\text{ext}}$ denoting the extended shape operator of a *smoothed* representation of the level set function w and n_σ denoting the associated smoothed normal. As a smoothing operator, at every $x \in \Omega$, one might consider the L^2 projection of w onto the space of quadratic polynomials in a ball $B_\sigma(x)$. Finally, spatial discretization can be achieved via finite element discretization of the computational domain (see Section 2). Figure 7 shows

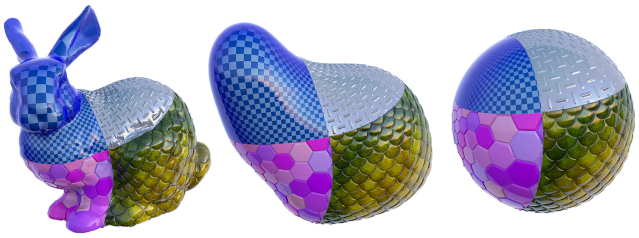

Fig. 9 Willmore flow on triangulated meshes when working in *curvature space*, see [54]. (Image courtesy of K. Crane.)

the application of the anisotropic diffusion method in the context of 3D medical imaging [48]. Furthermore, the method can be applied to surfaces with large amounts of geometric detail, when the implementation is based on hierarchical sparse narrow band structures [49].

Higher order geometric flows. Mean curvature flow, as discussed above, corresponds to a second order PDE. As such, it can only accommodate for positional boundary conditions, whereas prescribing tangential boundary conditions is not possible. This poses serious restrictions when attempting to construct surfaces with tangential continuity, e.g., for surface restoration or inpainting. This problem can be handled by working with higher order flows, such as *surface diffusion*, $\partial_t x - (\Delta h)n = 0$, see, e.g., Xu et al. [50] (and Schneider and Kobbelt [51] for a similar earlier approach).

A prominent fourth order flow arises from gradients of the Willmore energy $\frac{1}{2}\int_{\mathcal{S}}(h)^2\,\mathrm{d}a$. The resulting flow was, e.g., investigated by Yoshizawa and Belyaev [52] (using a discretization of the Euler–Lagrange equations), by Bobenko and Schröder (using a discrete geometric approach that preserves the Möbius invariance of the smooth formulation $\int_{\mathcal{S}}(\kappa_1-\kappa_2)^2\,\mathrm{d}a$), and in [53] (using a finite element formulation). All of these methods face the problem of time step restrictions. By working in *curvature space* instead of in position space, Crane et al. [54] recently proposed a formulation of Willmore flow for triangulated surfaces that allows for large time steps while preserving the quality of the input mesh, see Figure 9. Their approach can also be used for highly efficient surface fairing while preserving mesh quality.

4 Elastic Energy of Thin Shells

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

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

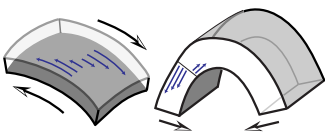

Fig. 10 Tangential (left) and bending (right) distortion.

For small thickness δ one mainly distinguishes (i) tangential distortion caused by in-layer (tangential) shear, compression, or expansion and (ii) (transversal) shear caused by normal bending. With respect to applications in geometry processing one is often 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 surfaces.

Tangential distortion energy. Let $\mathcal{S} \subset \mathbb{R}^3$ be a smooth, embedded surface. Consider a deformation $\phi : \mathbb{R}^3 \to \mathbb{R}^3$ (which we assume to be a diffeomorphism), and consider the resulting deformed surface $\phi(\mathcal{S})$. The mechanically relevant deformation tensor is the frame indifferent Cauchy–Green strain tensor $\mathbf{G}[\phi] = D\phi^T D\phi$, which, when restricted to the tangent space of $\mathcal{S}$, represents the pullback of the first fundamental from of the deformed surface to the undeformed configuration. This tangential component is defined by

$$\mathrm{I}_{\mathcal{S}}(\mathbf{G}^{\tan}[\phi](u), v) = \mathrm{I}_{\phi(\mathcal{S})}(D_{\mathcal{S}}\phi(u), D_{\mathcal{S}}\phi(v))$$

for all pairs (u, v) of tangent vectors at a point of $\mathcal{S}$.[4] When $x : \Omega \to \mathcal{S}$ denotes a local parameterization, let $\psi = \phi \circ x : \Omega \to \mathbb{R}^3$. Then the first fundamental forms of $\mathcal{S}$ and $\phi(\mathcal{S})$ are represented by $g^{\mathcal{S}} = Dx^T Dx$ and $g^{\phi(\mathcal{S})} = D\psi^T D\psi$, respectively (see Section 2), and $\mathbf{G}^{\tan}[\phi]$ takes the form $\mathbf{G}^{\Omega}[\phi] = (g^{\mathcal{S}})^{-1} g^{\phi(\mathcal{S})}$. 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 tensor and scales linearly in the thickness δ, i.e.,

$$\mathcal{W}_{\mathrm{mem}}[\phi] = \delta \int_{\mathcal{S}} W_{\mathrm{mem}}(\mathbf{G}^{\tan}[\phi]) \, \mathrm{d}a \,, \tag{8}$$

where the non-negative energy density W_{mem} acts on symmetric linear operators. For W_{mem} we require that (i) $W_{\mathrm{mem}}(\mathbb{1}) = 0$, (ii) $DW_{\mathrm{mem}} = 0$ at the identity matrix, and (iii) $D^2 W_{\mathrm{mem}}$ is positive definite at the identity matrix. These requirements correspond to the fact that if the shell is in a stress free configuration, then the deformation identity $\mathbb{1}$ is a minimizer of $\mathcal{W}_{\mathrm{mem}}$ and thus (i) $\mathcal{W}_{\mathrm{mem}}[\mathbb{1}] = 0$ and (ii) $d\mathcal{W}_{\mathrm{mem}}[\mathbb{1}] = 0$. Additionally, we assume (iii) that the energy is strictly convex (modulo rigid body motions) in a neighborhood of a minimizer. These assumptions capture most thin elastic materials [55].

Bending energy. Bending of thin shells is described in terms of the change in the variation of the normal on the surface. Hence, to quantify bending, one compares the shape operator $\mathbf{S}_{\phi(\mathcal{S})}$ on the deformed surface with the shape operator $\mathbf{S}_{\mathcal{S}}$ on the undeformed surface. To this end, one first pulls back $\mathbf{S}_{\phi(\mathcal{S})}$ to $\mathcal{S}$ by defining a linear operator $\mathbf{S}[\phi]$ via

$$\mathrm{I}_{\mathcal{S}}(\mathbf{S}[\phi](u), v) = \mathrm{I}_{\phi(\mathcal{S})}(\mathbf{S}_{\phi(\mathcal{S})}(D_{\mathcal{S}}\phi(u)), D_{\mathcal{S}}\phi(v)) \,,$$

for all pairs (u, v) of tangent vectors of $\mathcal{S}$. Then one defines the *relative* shape operator as

$$\mathbf{S}_{\mathrm{rel}}[\phi] = \mathbf{S}[\phi] - \mathbf{S}_{\mathcal{S}} \,.$$

Given a local parameterization $x : \Omega \to \mathcal{S}$, the relative shape operator is represented as $\mathbf{S}^{\Omega}_{\mathrm{rel}}[\phi] = g^{-1}(D(\phi \circ x)^T D(n^{\phi} \circ \phi \circ x) - Dx^T D(n \circ x))$. If ϕ is locally isometric, i.e., if

[4] Here $D_{\mathcal{S}} = d$ denotes exterior differentiation. We use $D_{\mathcal{S}}$ as it is more common in the mechanics community.

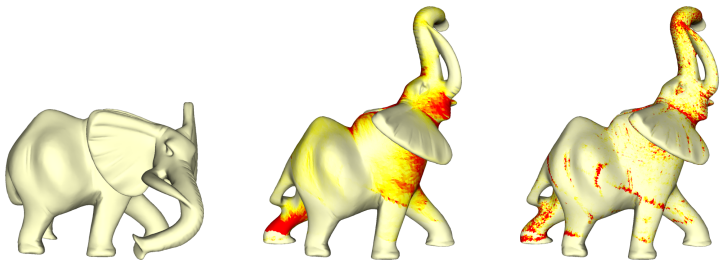

Fig. 11 Two predominant types of deformation for a surface model (left) in the deformed configuration using a discrete geometry approach on triangular meshes: tangential stretching of shearing (middle, 0 0.001) and normal bending (right, 0 0.1).

$\mathbf{G}^{\tan}[\phi] = \mathbb{1}$, then $\mathbf{S}_{\mathrm{rel}}[\phi]$ measures pure bending; otherwise, the relative shape operator also reflects tangential distortion. A suitable bending energy is then given by

$$\mathcal{W}_{\mathrm{bend}}[\phi] = \delta^3 \int_{\mathcal{S}} W_{\mathrm{bend}}(\mathbf{S}_{\mathrm{rel}}[\phi])\, \mathrm{d}a\,. \tag{9}$$

For W_{bend} we require that (i) $W_{\mathrm{bend}}(0) = 0$, (ii) $DW_{\mathrm{bend}} = 0$ at the zero matrix, and (iii) $D^2 W_{\mathrm{bend}}$ is positive definite at the zero matrix. 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 has thickness δ. A simple model for the energy density is $W_{\mathrm{bend}}(\mathbf{S}_{\mathrm{rel}}[\phi]) = |\mathbf{S}_{\mathrm{rel}}[\phi]|^2$, where $|A|$ denotes the Frobenius norm of the matrix A. Notice that $W_{\mathrm{bend}}(\mathbf{S}_{\mathrm{rel}}[\phi])$ takes into account the full change of the shape operators, not only the change of their traces (i.e., mean curvatures)—to the effect that changes of bending *directions* are accounted for appropriately.

Different surface representations. Observe that the above energy formulations apply whenever a formulation of first and second fundamental forms is available, e.g., for parametric surfaces, implicit surfaces, and triangle meshes.

5 Deformation and Physical Simulation of Thin Shells

In geometry processing, 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 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. 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 [56, 57]. Therefore, one may often assume that deformations nearly preserve the metric, i.e., are *nearly isometric*. Observe that if a deformation 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, and ∇ denotes covariant differentiation.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 $\mathrm{II}_{\phi(\mathcal{S})}(v, w)n^\phi = \mathrm{Hess}[\phi](v, w)$, 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 has been used in [53] to accelerate computations.

Linearized elasticity. While isometric deformations of thin plates lead to quadratic bending energies, this is no longer the case for thin shells that are not flat in their 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 small $\epsilon \in \mathbb{R}_+$. To further simplify the exposition, we additionally restrict to energy densities given by the Frobenius norm, i.e., $W_{\mathrm{mem}}(A) = |A - \mathbb{1}|^2$ and $W_{\mathrm{bend}}(A) = |A|^2$. Then the corresponding membrane energy is given by

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

where dots denote higher order terms with respect to ϵ and $(D_{\mathcal{S}} v)_{\mathrm{sym}}$ is the symmeterized derivative, i.e., $(D_{\mathcal{S}} v)_{\mathrm{sym}} = \mathbb{1}_{3\times 2}^T D_{\mathcal{S}} v + D_{\mathcal{S}} v^T \mathbb{1}_{3\times 2}$, 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

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

where $\mathrm{Hess}[v]$ denotes the component-wise Hessian of $v : \mathcal{S} \to \mathbb{R}^3$, 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 $\mathrm{Hess}[v] \cdot n$ into a symmetric 2×2 tensor field. The assumption of infinitesimally isometric deformations is motivated by considering *pure bending*, i.e., by disallowing any contribution of membrane terms in the bending energy.

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. In order to be able to work with standard solvers, the above energies are by some authors further modified to

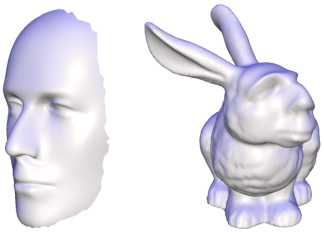

Fig. 12 Similar to mesh deformation, the linear editing metaphor can be used for coating transfer (e.g., from mannequin to bunny) [58]. (Image courtesy of O. Sorkine–Hornung.)

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

where $\|\cdot\|$ denotes the usual Euclidean norm. Notice that different from $|(D_{\mathcal{S}} v)_{\mathrm{sym}}|$, which does not account for infinitesimal rotations, $\|D_{\mathcal{S}} v\|$ is the full norm and therefore also incorporates rotations and thus contains 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 $\triangle_{\mathcal{S}} v = \mathrm{tr}\,\mathrm{Hess}[v]$, does not account for principal bending directions. From an implementation point of view, however, 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 [59] for a discussion of benefits and drawbacks of the linearized approach in computer graphics applications.

Restoring rotation invariance. A major limitation of linear elasticity is that it comes at the price of *losing rotation invariance*. Several approaches in the literature have attempted to overcome this limitation. 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 [60]. 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

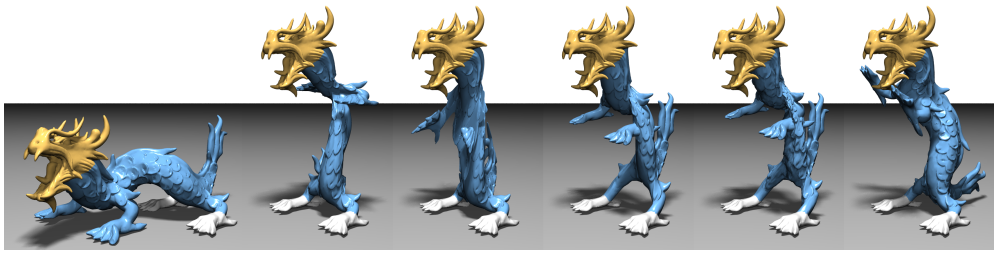

Fig. 13 Deformation of the dragon (left) by fixing its hind feet and moving its head upwards in a single step. The result of PriMo [61] (right) compared to linear deformation methods (second through fifth picture) that tend to yield counter-intuitive results. (Image courtesy of M. Botsch.)

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 [59] for a discussion of pros and cons for various methods that employ this methodology.

Another set of approaches builds on giving up linearity in order to restore rotation invariance. One example of a popular, efficient, and robust method for nonlinear surface deformation and shape matching is PriMo [61]. This approach is based on replacing the triangles of a polyhedral surfaces by thin prisms. During a deformation, these prisms are required to stay rigid, while nonlinear elastic forces are acting between neighboring prisms to account for bending, twisting, and stretching of the surface. This requires to solve for optimal rigid motions under the constraints prescribed by the user. Another nonlinear method for restoring rotation invariance was suggested by Sorkine and Alexa in their work on *as-rigid-as-possible* (ARAP) deformations [62]. 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. Unfortunately, neither PriMo nor ARAP appear to have a continuum mechanics analogue.

Volume deformations. Another approach that is prevalent in the literature is to view surfaces as boundaries of 3D volumes. In this setting, deformations are 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 $\phi : M \subset \mathbb{R}^3 \to \mathbb{R}^3$ to the rotation group $\mathrm{SO}(3)$, by minimizing the energy

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

under user-specified constraints. This approach is based on the observation that, apart from *globally* rigid motions, locally varying rotations will unavoidably induce deviations from isometry. To be precise, suppose without loss of generality (by factoring out globally rigid motions) that a volume deformation ϕ fixes the center of mass and induces no global rotation,

Fig. 14 Snapshots from the simulation of a billowing flag. Despite its economy of cost, the proposed *isometric bending model* in [53] achieves qualitatively the same dynamics as popular nonlinear models.

i.e., $\int_M \phi \,\mathrm{d}x = \int_M \mathrm{curl}(\phi)\,\mathrm{d}x = 0$. Then *geometric rigidity* [63] implies that there exists a constant C depending on M (but not on ϕ) such that

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

This inequality reflects the fact that if an infinitesimal piece of material is rotated with respect to a "neighboring" infinitesimal piece, then such a rotation induces local stretching, compression, or shearing of the material. Notice the conceptual similarity of this observation to the approaches employed by PriMo and ARAP discussed above.

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. [64] 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—with one important difference: 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 artifacts. The early graphics literature has focused on *efficiency*—and, to achieve this, has sometimes willingly been 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, i.e., $F_{\mathrm{elastic}}[\phi] = -\nabla\mathcal{W}[\phi] = -\nabla\mathcal{W}_{\mathrm{mem}}[\phi] - \nabla\mathcal{W}_{\mathrm{bend}}[\phi]$.

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 [65, 66]. In the Rayleigh view, the damping force, F_{damp}, is proportional to velocity, $v(t) = \dot{\phi}(t)$:

$$F_{\mathrm{damp}}[v] = -K_{\mathrm{d}} v \,, \quad \text{where} \quad K_{\mathrm{d}} = \alpha_1 \mathbf{M} + \alpha_2 \mathrm{Hess}[\mathcal{W}] \,. \tag{10}$$

Here, the linear operator K_d is written as a linear combination[5] of two tensors: the mass matrix $\mathbf{M}$ 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. [67].

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:

$$\begin{pmatrix} Id & 0 \\ 0 & \rho\mathbf{M} \end{pmatrix} \begin{pmatrix} \dot{\phi}(t) \\ \dot{v}(t) \end{pmatrix} = \begin{pmatrix} v(t) \\ F_{\mathrm{elastic}}[\phi(t)] + F_{\mathrm{damp}}[v(t)] + F[\phi(t), v(t))] \end{pmatrix} , \tag{11}$$

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*, $\rho\mathbf{M}$, is given by the product of mass surface density and the geometric mass matrix. In order to accelerate force computations, it is often desirable to have *explicit* representations of energy gradients and energy Hessians instead of numerically deriving these quantities on the fly by automatic differentiation.

Time discretization of (11) is a well-studied problem (see [68] and references therein); methods may be classified as explicit, implicit, or mixed implicit-explicit.[6] *Geometric (or variational) integrators* [69] have been advocated due to their *structure preservation*, e.g., guaranteed preservation of momenta and near preservation of energy even for large simulation times.

6 Matching of Thin Shells

The feature aware matching of two given surfaces is one of the fundamental tasks in 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

[5] 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.

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

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. There exists a vast amount of literature on matching triangulated surfaces. For example, starting from the notion of Gromov–Hausdorff distances, Bronstein et al. [70] have developed an efficient algorithm for the matching of triangular surfaces. Their method also applies to the case of partial correspondence. A topologically robust variant of this approach based on diffusion distances was proposed in [71]. The observation that robustness to noise and changes of topology can be significantly increased by a multiscale approach has also been utilized by Sun, Ovsjanikov, and Guibas in [72], 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. Physically, one might consider surface matching as the pressing of a given surface $\mathcal{S}_A$ into the mould of another surface $\mathcal{S}_B$. Given local parameterizations $x_A : \Omega_A \to \mathcal{S}_A$ and $x_B : \Omega_B \to \mathcal{S}_B$ of the two surfaces, the matching problem can be phrased in terms of a functional that yields an optimal mapping $\phi^\Omega : \Omega_A \to \Omega_B$.[7] The attendant matching functional can be composed as a sum of terms that penalize membrane and bending distortions induced by ϕ^Ω. Additionally, besides the tangential distortion energy $\mathcal{W}_{\mathrm{mem}}$, which can be regarded as a regularization energy ensuring smoothness of the matching deformation, and the bending energy $\mathcal{W}_{\mathrm{bend}}$, which attempts to match regions of equal curviness, such as corresponding creases, one might want to align preselected feature regions, such as the eyes on facial surfaces via a third energy term $\mathcal{W}_{\mathrm{feature}}$. Defining $\phi = x_B \circ \phi^\Omega \circ x_A^{-1}$, one obtains the tangential Cauchy–Green tensor $\mathbf{G}^{\tan}[\phi]$ and the relative shape operator $\mathbf{S}_{\mathrm{rel}}[\phi]$ as described above, from which one assembles $\mathcal{W}_{\mathrm{mem}}$ and $\mathcal{W}_{\mathrm{bend}}$, respectively. In particular, for the membrane energy density one can use $W_{\mathrm{mem}}(A) = \hat{W}(a,b) = \alpha_l a + \alpha_a\big(b + (1 + \frac{\alpha_l}{\alpha_a})\, b^{-1}\big) - 2\alpha_a - 3\alpha_l$, where $a = a(A) = \operatorname{tr} A$ accounts for length distortion and $b = b(A) = \det A$ accounts for area expansion with b and

Fig. 15 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 map can be mapped between surfaces (center). Unmatched regions are depicted in black.

[7] For the case of partial matching, one allows for $\phi^\Omega(\Omega_A) \neq \Omega_B$.

area compression with b^{-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 [73] was rigorously derived in [74] from a set of natural axioms for measuring the distortion of a single parameterization. If one is only interested in the comparison of the mean curvatures of the two surfaces, one can use the bending energy density $W_{\mathrm{bend}}(A) = (\operatorname{tr} A)^2$. Finally, if we denote by $\mathcal{F}_A \subset \Omega_A$ and $\mathcal{F}_B \subset \Omega_B$ feature sets in the parameter domains of the two surfaces, then

$$\mathcal{W}_{\mathrm{feature}}[\phi^\Omega] = \int_{\Omega_A} \left((\chi_{\mathcal{F}_B} \circ \phi^\Omega)(1 - \chi_{\mathcal{F}_A}) + (1 - \chi_{\mathcal{F}_B} \circ \phi^\Omega)\chi_{\mathcal{F}_A} \right) \sqrt{\det g_A}\, d\xi$$

measures the symmetric difference on the surface A of the feature set on surface A and the pullback of the feature set on the surface B. Here χ_Ω denotes the characteristic function of a set Ω. The resulting combined energy $\mathcal{W}[\phi^\Omega]$ is then given by an appropriately weighted sum of $\mathcal{W}_{\mathrm{mem}}[\phi^\Omega]$, $\mathcal{W}_{\mathrm{bend}}[\phi^\Omega]$, and $\mathcal{W}_{\mathrm{feature}}[\phi^\Omega]$. Fig. 15 shows an application of surface matching.

Matching implicit surfaces. Given two surfaces $\mathcal{S}_A$ and $\mathcal{S}_B$, one can use an implicit approach by representing them by their respective signed distance functions d_A and d_B. In this case one obtains for the extended shape operator (see above) $\mathbf{S}_X^{\mathrm{ext}} = Dn_X P_X = D^2 \mathrm{d}_X$ for $X \in \{A, B\}$, where $P_X = \mathbb{1} - \nabla \mathrm{d}_X \otimes \nabla \mathrm{d}_X$ is the projection to a level set of the respective distance function. One choice of a robust and efficient algorithm for the conversion of a general implicit surface representation into a signed distance function is the fast marching method [75]. Since we are interested in matching the two surfaces $\mathcal{S}_A$ and $\mathcal{S}_B$, a narrow band approach enables "blending out" distant implicit surfaces $[\mathrm{d}_A = c]$ and $[\mathrm{d}_B = c]$ for $|c| \geq \epsilon$. Indeed, 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 (and the fact that $\|\nabla \mathrm{d}_A\| = 1$) for a deformation $\phi : \mathbb{R}^3 \to \mathbb{R}^3$, one obtains the tangential distortion energy on the narrow band $[-\sigma \leq \mathrm{d}_A \leq \sigma]$ as given by

$$\mathcal{W}_{\mathrm{mem}}[\phi] = \int_{\mathbb{R}^3} (\eta_\sigma \circ \mathrm{d}_A)\, W_{\mathrm{mem}}(\mathbf{G}^{\tan}[\phi] + \nabla \mathrm{d}_A \otimes \nabla \mathrm{d}_A)\, \mathrm{d}x\,,$$

where $\mathbf{G}^{\tan}[\phi]$ is the above defined tangential Cauchy–Green strain tensor, trivially extended to a 3×3 tensor by requiring that $\mathbf{G}^{\tan}[\phi] n_A = 0$. As a first approach for the bending energy one gets (see [76])

$$\mathcal{W}_{\mathrm{bend}}[\phi] = \int_{\mathbb{R}^3} (\eta_\sigma \circ \mathrm{d}_A) |P_A D\phi^T ((D^2 \mathrm{d}_B) \circ \phi)\, D\phi P_A - D^2 \mathrm{d}_A|^2\, \mathrm{d}x\,.$$

Unfortunately, the resulting energy is not lower semicontinuous. To remedy this deficiency, we modify the bending term by first modifying the definition of the tangential Jacobian by setting $D_A^B \phi = (P_B \circ \phi) D\phi P_A$ and then defining

$$\mathcal{W}_{\mathrm{bend}}[\phi] = \int_{\mathbb{R}^3} (\eta_\sigma \circ \mathrm{d}_A) W \left(((\mathbf{S}_B^{\mathrm{ext}} + \beta \mathbb{1})^{\frac{1}{2}} \circ \phi)\, D_A^B \phi\, (\mathbf{S}_A^{\mathrm{ext}} + \beta \mathbb{1})^{-\frac{1}{2}} + \Lambda \right) \mathrm{d}x\,,$$

where $\Lambda = \nabla \mathrm{d}_B \otimes \nabla \mathrm{d}_A$. This modification indeed yields a lower semicontinuous functional for large enough β, see [77]. Here, the integrand W is chosen similarly to the integrand

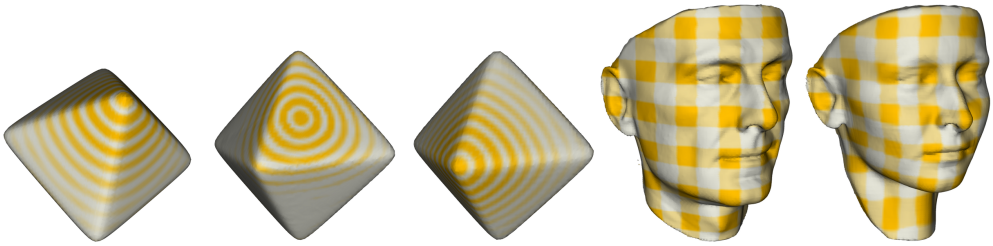

Fig. 16 Matching of implicit surfaces based on the simplified thin shell model [76] for two octahedrons (left) (the second image shows a matching without bending energy) and two different faces (right). Thereby the texture is deformed together with the template surface and shows the proper matching of geometric features.

of *membrane* energy. Besides the membrane and bending energies on the narrow band, one requires a penalty functional,

$$\mathcal{W}_{\text{penalty}}[\phi] = \frac{1}{\epsilon} \int_{\mathbb{R}^3} (\eta_\sigma \circ \mathrm{d}_A)(\mathrm{d}_A - \mathrm{d}_B \circ \phi)^2 \, \mathrm{d}x \,,$$

in order 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 to ensure in addition that the deformation ϕ is regular and injective also outside the narrow band in the region of computation Ω around $\mathcal{S}_A$. To this end, we incorporate a regularization energy,

$$\mathcal{W}_{\text{reg}}[\phi] = \int_\Omega (1 - \eta_\sigma \circ \mathrm{d}_A) W_{\text{reg}}(C[\phi]) \, \mathrm{d}x \,,$$

where $C[\phi] = D\phi(x)^T \, D\phi(x)$ is the three-dimensional Cauchy–Green strain tensor and $W_{\text{reg}}(C) = \frac{\mu}{4}\mathrm{tr}(C)^2 + (\lambda - 2\mu)\det(C)^{\frac{1}{2}} - \frac{\lambda-\mu}{2}\log\det(C)$ for Lamé-Navier coefficients λ and μ [78]. A weighted sum of the energy terms $\mathcal{W}_{\text{mem}}$, $\mathcal{W}_{\text{bend}}$, $\mathcal{W}_{\text{penalty}}$, and $\mathcal{W}_{\text{reg}}$, yields the total elastic energy. Similar to the parametric case, a feature matching energy can additionally be taken into account.

7 Spectral and Modal Methods

Building on a general paradigm of global analysis, it can be advantageous to switch perspective and describe surfaces through function spaces (such as Sobolev spaces) and differential operators (such as the Laplacian) acting between them. This *functional perspective* has recently been advocated for various geometry processing applications by Ovsjanikov, Ben-Chen, and others [79–81]. For (formally) self-adjoint differential operators (such as the Laplace–Beltrami operator), one can consider the spectrum and the eigenfunctions as descriptors of the surface, which allows for applying classical signal processing approaches to irregular triangle meshes. Indeed, the eigenfunctions of the Laplacian constitute a basis that provides the analogue of the Fourier basis in the planar case. Geometry processing applications that utilize the spectral viewpoint include mesh compression [82,83], mesh parameterization [84], mesh fingerprinting [85], mesh segmentation, mesh registration [86], and so forth.

Fig. 17 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.)

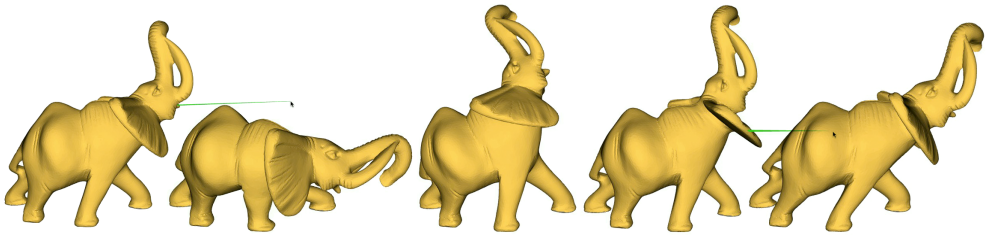

Fig. 18 Results of an interactive animation of an elephant model using modal analysis to build a reduced basis of forces [88]. (Image courtesy of K. Hildebrandt.)

Another operator that is relevant from the perspective of shell deformations is the *Hessian* of the elastic energy. If the Hessian of elastic energy is considered at the (undeformed) rest state 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 in the normal direction of a surface that describes an (infinitesimal) oscillation of the elastic shell with frequency λ_i, where low eigenvalues correspond to a low degree of stiffness and thus to a physically preferred mode of shell variability.

The spectrum and eigenvalues of this energy Hessian provide the so-called *modal basis*, which is in general different from the spectral basis, and that is adapted to elastic properties of a given surface. Hildebrandt et al. [87, 88] showed how to use this basis for the intuitive modeling of surfaces and for accelerating physical simulations. Indeed, for accelerating physical simulations, a starting point is to observe that if one considers Rayleigh damping, then the equations of motion decouple (as ordinary differential equations) in the modal basis.

As a significantly simplified, linear model one considers solely normal variations un 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 17. Here, a projected inverse vector iteration has been applied to the discrete weak form (6) in the Loop subdivision finite element approach.

8 From shapes to the space of shapes

Up to this point, we have discussed computational tools to process single geometries. From a more global perspective one might want to study surfaces as points in a *space of shapes*. 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

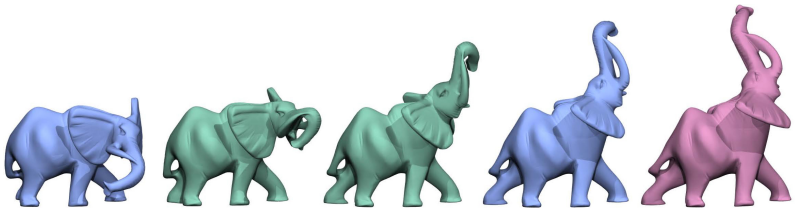

Fig. 19 Two intermediate surface (green) from a geodesic between two poses (blue) of an elephant model and a extrapolation via the exponential map (magenta) using the approach by Kilian *et al.* [90]. (Image courtesy of M. Kilian.)

morphing and modeling [89, 90], in computational anatomy [91, 92], as well as shape statistics [93, 94]. Studying shape space from the point of view of Riemannian geometry enables transfer of important concepts from classical geometry to these (usually) infinite-dimensional spaces. Examples with a fully developed geometric theory include spaces of planar curves with curvature based metrics [95], elasticity based metrics [96, 97], or Sobolev-type metrics [98]. In a geometrically motivated approach, Kilian et al. consider geodesics between consistently triangulated surfaces (see [90] and Fig. 19), where the underlying Riemannian metric measures the stretching of triangle edges. While this metric is invariant to rigid body motions, the lack of a bending term leaves a non-trivial kernel of the metric tensor, including all isometric deformations of the triangular mesh. To avoid the resulting unphysical wrinkling effects, a supplementary (nonphysical) regularization was incorporated by Killian and co-workers. Instead, one may use the regularizing effect of bending energy—and stay entirely in the realm of a physically sound theory—as demonstrated in [99, 100].

Time-continuous setting. The central observation, both for smooth surfaces and triangle meshes, is that the *Hessian of elastic energy*[8] when evaluated at the undeformed surface is *positive semidefinite*, and its kernel consists of infinitesimal rigid transformations only, see [100]. Thus, modulo rigid transformations of surfaces, this Hessian provides a notion of a Riemannian metric on the space of shapes, which we denote by $\mathbf{g}_{\mathcal{S}}(\cdot,\cdot)$ at a point $\mathcal{S}$ in shape space. Notice that a tangent vector v at a point $\mathcal{S}$ in shape space corresponds to a vector field $v : \mathcal{S} \to \mathbb{R}^3$. Geodesics that are associated with this Riemannian metric have a physical interpretation. Indeed, viewing surfaces as thin *viscous* materials[9], a geodesic is the deformation path between two points in the space of shells that has the *least energy dissipation*. To see this, consider a family $(\phi_t)_{0\le t\le 1}$ of diffeomorphisms of $\mathbb{R}^3$ with $\phi_0 = \mathbb{1}$. This family generates a deformation path $\phi_t(\mathcal{S})$, and if shells were purely *viscous*, then along this path energy would be dissipated. According to Rayleigh's analogy that derives a viscous formulation from an elastic one by replacing elastic strain by strain *rates* [101], the total energy dissipation is

$$\mathcal{E}[(\phi_t)_{0\le t\le 1}] = \int_0^1 \operatorname{Hess} \mathcal{W}_t[\mathbb{1}](v_t, v_t)\,\mathrm{d}t = \int_0^1 \mathbf{g}_{\mathcal{S}_t}(v_t, v_t)\,\mathrm{d}t\,. \tag{12}$$

Here $v_t = \dot{\phi}_t \circ \phi_t^{-1}$ is the Eulerian velocity field associated with the deformation, and $\mathcal{W}_t$ is the elastic energy associated with the shell $\mathcal{S}_t = \phi_t(\mathcal{S})$, where for each fixed t, the shell

[8] As before, elastic energy is a sum of membrane *and* bending contributions.

[9] Although the theory initially requires finite thickness, we restrict to the midsurface to simplify the expositions.

Fig. 20 Discrete 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).

$\mathcal{S}_t$ is the undeformed rest state of $\mathcal{W}_t$. In Riemannian geometry, the right hand side of (12) is known as *path energy*, and its minimizers are constant speed geodesics, thus establishing a rigorous notion of distance between surfaces. The resulting Riemannian structure on the space of shells not only gives rise to the notion of shortest paths. Additionally, it allows for defining (i) the logarithm (mapping shortest paths between two shapes into an initial velocity vector which "shoots" to the target shape), (ii) the exponential (mapping an initial shooting direction into a particular shape via a shortest path) and (iii) the notion of parallel transport. We discuss the corresponding constructions below.

Time-discrete setting. For computational feasibility, one requires a time-discrete analogue of the time-continuous setting outlined above. Here, we offer a brief outline of such a discrete theory. For details we refer to [99, 100, 102, 103]. Consider the time-discrete family $(\mathcal{S}_k)_{k=0,\dots,K}$ of shells given by $\mathcal{S}_k = \phi_{k\tau}(\mathcal{S})$, where $\tau = \frac{1}{K}$ denotes the discrete time step. A straightforward Taylor expansion shows that one can approximate, up to second order in τ, the path energy (12) by

$$\mathcal{E}[\mathcal{S}_0, \dots, \mathcal{S}_K] = K \sum_{k=1}^{K} \mathcal{W}[\mathcal{S}_{k-1}, \mathcal{S}_k] \,, \tag{13}$$

where $\mathcal{W}[\mathcal{S}_{k-1}, \mathcal{S}_k]$ denotes the elastic energy required for deforming $\mathcal{S}_{k-1}$ into $\mathcal{S}_k$. A discrete geodesic path is now defined as a minimizer of the discrete path energy $\mathcal{E}[\mathcal{S}_0, \dots, \mathcal{S}_K]$ with given boundary conditions $\mathcal{S}_0 = \mathcal{S}_A$ and $\mathcal{S}_K = \mathcal{S}_B$. Concerning the spatial discretization for triangulated meshes, one might use the notions of discrete first and second fundamental forms discussed in Section 2 for defining elastic energies. Figure 20 shows a discrete geodesic together with the different components of the underlying local dissipation rate.

In what follows, we 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 $\mathcal{W}^h[\bar{X}_{k-1}, \bar{X}_k]$ and $\mathcal{E}^h[\bar{X}_0, \dots, \bar{X}_K]$, respectively. If $(X = \bar{X}_0, \dots, \bar{X}_K = Y)$ is a discrete geodesic, then the initial displacement $\bar{X}_1 - \bar{X}_0$

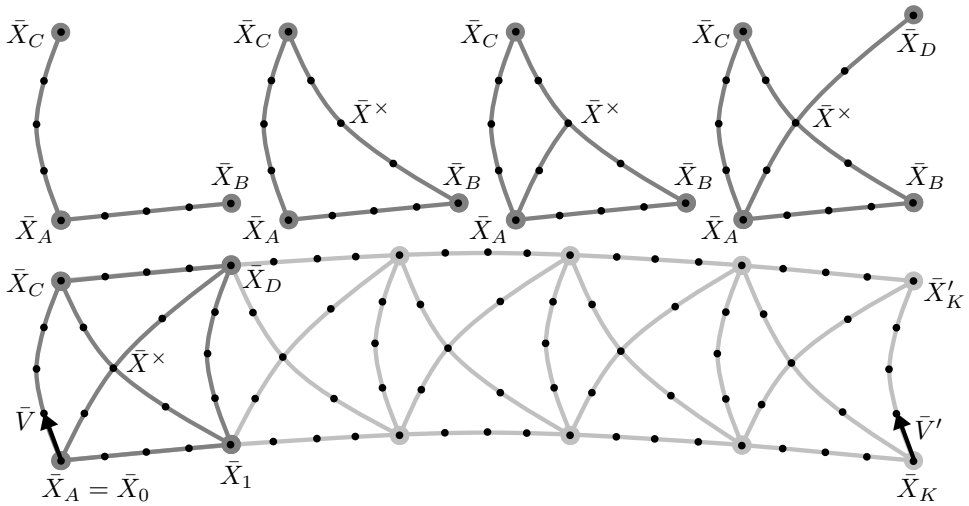

Fig. 22 Top: Stepwise construction of a (discrete) geodesic parallelogram (in the discrete case $K = 4$ such that each discrete geodesic consists of five shapes indicated by the black dots). Bottom: Iterative construction for the time-discrete parallel transport of $\bar{V}$ along $(\bar{X}_0, \ldots, \bar{X}_K)$ to $\bar{V}'$.

scaled with $\frac{1}{\tau}$ can be considered as a time-discrete counterpart of the time-continuous geometric logarithm $\log_{\bar{X}} \bar{Y}$, which gives rise to the notation $\frac{1}{K}\mathrm{Log}_{\bar{X}}\bar{Y} = \bar{X}_1 - \bar{X}_0$. The time-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 time-discrete exponential map $\mathrm{Exp}_{\bar{X}}^K \bar{V}$ (with K time steps) as that position $\bar{Y}$, such that the initial displacement of the resulting discrete geodesic $[\bar{X} = \bar{X}_0, \ldots, \bar{X}_K = \bar{Y}]$ between $\bar{X}$ and $\bar{Y}$ equals $\bar{V}$, i.e., $\bar{V} = \bar{X}_1 - \bar{X}_0$. Figure 21 shows results of the discrete exponential map. Finally, we outline how to proceed with respect to a discrete parallel transport. To define a time-discrete parallel transport $\mathbf{P}_{\bar{X}_0,\ldots,\bar{X}_K}\bar{V}$ of a displacement $\bar{V}$ along a discrete curve $(\bar{X}_0, \ldots, \bar{X}_K)$ one can use a time-discrete counterpart of a Riemannian parallelogram construction, also known as *Schild's ladder* in gravitational physics [104, 105]: Given three shells, represented as $\bar{X}_A$, $\bar{X}_B$, and $\bar{X}_C$, respectively, one can construct a fourth shell, represented by $\bar{X}_D$, such that the geodesic paths between the four shells form a geodesic parallelogram, i.e., the geodesics $\bar{X}_C \leftrightarrow \bar{X}_B$ and $\bar{X}_A \leftrightarrow \bar{X}_D$ share the same midpoint $\bar{X}^\times$ (see Fig. 22). To compute $\bar{X}_D$ we proceed as follows: (i) compute the geodesic from $\bar{X}_C$ to $\bar{X}_B$ via interpolation; (ii) take its midpoint $\bar{X}^\times$, (iii) compute the geodesic from $\bar{X}_A$ to $\bar{X}^\times$, (iv) compute $\bar{X}_D$ through extrapolation from $\bar{X}^\times$ with initial direction $\bar{V}^\times$, where $\bar{V}^\times$ is the final velocity of the geodesic in step (iii) at its endpoint $\bar{X}^\times$. A sequence of K such parallelograms then transports the initial velocity $\bar{V}$ of $\bar{X}_A \leftrightarrow \bar{X}_C$ along the path from $\bar{X}_A = \bar{X}_0$ to $\bar{X}_K$ to the velocity $\bar{V}'$ of $\bar{X}_K \leftrightarrow \bar{X}'_K$ (see also Fig. 22).

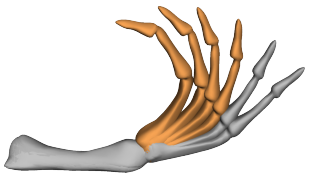

Fig. 21 Two positions of a finger as initial shapes (grey) and several steps of the discrete exponential map.

9 Conclusion

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 biased personal perspective. We attempted to cover various geometric representations including parametric and triangulated surfaces, level sets, and subdivision surfaces. Furthermore, we sketched the increasing interaction of mathematics, computer graphics, and mechanics. This interplay is characterized by fundamental and 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 there exist discrete (local) minimizers in the vicinity of a continuous (local) minimizer. It might also be rewarding to study the curvature of shape space based in the above introduced Riemannian metric. 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 by now, incorporating the structure of the space of shapes poses many challenging open problems.

Acknowledgements

We are grateful to the anonymous reviewers for their suggestions and their careful reading of our manuscript. M.W. was partially supported by the BMBF collaborative project MUSIKA.

References

[1] M. S. Floater and K. Hormann, Surface Parameterization: A Tutorial and Survey, in: Advances in Multiresolution for Geometric Modeling, edited by N. A. Dodgson, M. S. Floater, and M. A. Sabin (Springer, 2005), pp. 157–186.
[2] Y. Lipman, ACM Transactions on Graphics **31**(4), 108 (2012).
[3] R. MacNeal, The solution of partial differential equations by means of electrical networks, 1949, PhD thesis, California Institute of Technology.
[4] U. Pinkall and K. Polthier, Experim. Math. **2**, 15–36 (1993).
[5] G. Dziuk, Finite elements for the Beltrami operator on arbitrary surfaces, in: Partial Differential Equations and Calculus of Variations, edited by S. Hildebrandt and R. Leis, Lecture Notes in Mathematics Vol. 1357 (Springer, 1988), pp. 142–155.
[6] K. Hildebrandt, K. Polthier, and M. Wardetzky, Geometricae Dedicata **123**, 89–112 (2006).
[7] E. Grinspun, A. N. Hirani, M. Desbrun, and P. Schröder, Discrete shells, in: Sympos. Comput. Anim., (2003), pp. 62–67.
[8] J. M. Morvan, Generalized Curvatures, Geometry and Computing (Springer, 2008).
[9] D. Cohen-Steiner and J. M. Morvan, Sympos. Comput. Geom. pp. 312–321 (2003).
[10] D. Cohen-Steiner and J. M. Morvan, J. Differential Geom. **73**(3), 363–394 (2006).
[11] K. Hildebrandt and K. Polthier, Computer Graphics Forum **23**(3), 391–400 (2004).
[12] E. Grinspun, Y. Gingold, J. Reisman, and D. Zorin, Eurographics (Computer Graphics Forum) **25** (2006).
[13] E. Catmull and J. Clark, Computer Aided Design **10**, 350–355 (1978).
[14] C. Loop, Smooth spline surfaces over irregular meshes, in: Proceedings of ACM SIGGRAPH, (1994), pp. 303–310.
[15] C. Mandal, H. Qin, and B. Vemuri, Computer-Aided Design **32**(8-9), 479–497 (2000).
[16] F. Cirak, M. Ortiz, and P. Schröder, International Journal for Numerical Methods in Engineering **47(12)**, 2039–72 (2000).

[17] E. Grinspun, P. Krysl, and P. Schröder, CHARMS: A Simple Franmework for Adaptive Simulation, in: Proceedings of ACM SIGGRAPH, (2002).
[18] G. Dziuk and C. M. Elliott, Math. Comp. **82**(281), 1–24 (2013).
[19] J. A. Sethian, Level Set Methods: Evolving Interfaces in Geometry, Fluid Mechanics, Computer Vision and Materials Sciences (Cambridge Univ. Press, 1996).
[20] S. Osher and R. P. Fedkiw, Journal of Computational Physics **169**(2), 463 – 502 (2001).
[21] B. Houston, M. Nielsen, C. Batty, O. Nilsson, and K. Museth, ACM Transactions on Graphics **25**(1), 1–24 (2006).
[22] J. D. Boissonnat, ACM Transactions on Graphics **3**, 266–286 (1984).
[23] H. Edelsbrunner, Geometry and topology for mesh generation (Cambridge University Press, 2001).
[24] N. Amenta, S. Choi, T. Dey, and N. Leekha, A simple algorithm for homeomorphic surface reconstruction, in: ACM Sympos. Comput. Geom., (2000), pp. 213–222.
[25] G. Turk and J. O'Brien, ACM Transactions on Graphics **21**, 855–873 (2002).
[26] A. Nealen, An as-short-as-possible introduction to the moving least squares method for scattered data approximation and interpolation, 2004, Technical Report.
[27] M. Kazhdan, M. Bolitho, and H. Hoppe, Poisson surface reconstruction, in: Proceedings of the Symposium of Geometry Processing, (2006), pp. 61–70.
[28] M. Belkin, J. Sun, and Y. Wang, Constructing Laplace operator from point clouds in r^d, SODA '09 (2009), pp. 1031–1040.
[29] M. Desbrun, A. Hirani, M. Leok, and J. E. Marsden(2005), `arXiv:math.DG/0508341`.
[30] M. Alexa and M. Wardetzky, ACM Transaction on Graphics **30**(4), 102:1–102:10 (2011).
[31] P. Mullen, P. Memari, F. deGoes, and M. Desbrun, ACM Transactions on Graphics **30**(4), 103:1–103:12 (2011).
[32] F. deGoes, P. Memari, P. Mullen, and M. Desbrun, ACM Transactions on Graphics **33**(3), 28:1–28:13 (2014).
[33] D. N. Arnold, R. S. Falk, and R. Winther, Acta Numerica **15**, 1–155 (2006).
[34] A. I. Bobenko and R. Seiler, Discrete Integrable Geometry and Physics (Oxford University Press, 1999).
[35] A. I. Bobenko, J. M. Sullivan, P. Schröder, and G. Ziegler, Discrete Differential Geometry, Oberwolfach Seminars (Birkhäuser Basel, 2008).
[36] A. I. Bobenko and Y. B. Suris, Discrete Differential Geometry. Integrable Structure (American Mathematical Society, 2008).
[37] M. Wardetzky, S. Mathur, F. Kälberer, and E. Grinspun, Discrete Laplace operators: No free lunch, in: Proceedings of the Symposium on Geometry Processing, (2007), pp. 33–37.
[38] K. Brakke, Exper. Math. **1**(2), 141–165 (1992).
[39] M. Desbrun, M. Meyer, P. Schröder, and A. Barr, Implicit fairing of irregular meshes using diffusion and curvature flow, in: Proceedings of ACM SIGGRAPH, (1999), pp. 317–324.
[40] G. Dziuk, Numer. Math. **58**, 603–611 (1991).
[41] G. Huisken, Journal für die reine und angewandte Mathematik **382**, 35–48 (1987).
[42] P. Perona and J. Malik, IEEE Transactions on Pattern Analysis and Machine Intelligence **12**(7), 629–639 (1990).
[43] B. Kawohl and N. Kutev, Math. Ann. **311**(1), 107–123 (1998).
[44] F. Catté, P. L. Lions, J. M. Morel, and T. Coll, SIAM J. Numer. Anal. **29**(1), 182–193 (1992).
[45] J. Weickert, Anisotropic diffusion in image processing (Teubner, 1998).
[46] K. Hildebrandt and K. Polthier, Computer Graphics Forum **23**(3), 391–400 (2004).
[47] C. Bajaj and G. Xu, ACM Transactions on Graphics **22**(1), 4–32 (2003).
[48] T. Preußer and M. Rumpf, SIAM Journal on Applied Mathematics **62**(5), 1772–1793 (2002).
[49] O. Nemitz, M. B. Nielsen, M. Rumpf, and R. Whitaker, SIAM Journal on Scientific Computing **31**(3), 2258–2281 (2009).
[50] G. Xu, Q. Pan, and C. L. Bajaj, Computer Aided Geometric Design **23**(2), 125–145 (2006).
[51] R. Schneider and L. Kobbelt, Computer Aided Geometric Design **18**(4), 359–379 (2001).

[52] S. Yoshizawa and A. G. Belyaev, Proceedings of the 2nd Biennial International Conference on Geometric Modeling and Processing pp. 119–123 (2002).
[53] M. Wardetzky, M. Bergou, D. Harmon, D. Zorin, and E. Grinspun, Computer Aided Geometric Design **24**, 499–518 (2007).
[54] K. Crane, U. Pinkall, and P. Schröder, ACM Transactions on Graphics **32**(4), 61:1–61:10 (2013).
[55] P. G. Ciarlet, Mathematical Elasticity, Vol III: Theory of Shells (North-Holland, 2000).
[56] T. Baraff and A. Witkin, Large steps in cloth simulation, in: Proceedings of ACM SIGGRAPH, (1998), pp. 43–54.
[57] D. Terzopoulos, J. Platt, A. Barr, and K. Fleischer, Elastically deformable models, in: Proceedings of ACM SIGGRAPH, (1987), pp. 205–214.
[58] O. Sorkine, D. Cohen-Or, Y. Lipman, M. Alexa, C. Rössl, and H. P. Seidel, Laplacian surface editing, in: Symposium on Geometry Processing, (ACM Press, 2004), pp. 179–188.
[59] M. Botsch and O. Sorkine, IEEE Transactions on Visualization and Computer Graphics **14**, 213–230 (2008).
[60] Y. Lipman, O. Sorkine, D. Levin, and D. Cohen-Or, ACM Transactions on Graphics **24**(3), 479–487 (2005).
[61] M. Botsch, M. Pauly, M. Gross, and L. Kobbelt, PriMo: Coupled prisms for intuitive surface modeling, in: Proceedings of the Symposium on Geometry Processing, (2006), pp. 11–20.
[62] O. Sorkine and M. Alexa, As-rigid-as-possible surface modeling, in: Symposium on Geometry Processing, (2007), pp. 109–116.
[63] G. Friesecke, R. James, and S. Müller, Comm.Pure Appl.Math. **55**(11), 1461–1506 (2002).
[64] I. Chao, U. Pinkall, P. Sanan, and P. Schröder, A simple geometric model for elastic deformations, in: Proceedings of ACM SIGGRAPH, (2010), pp. 38:1–38:6.
[65] T. J. R. Hughes, Finite Element Method - Linear Static and Dynamic Finite Element Analysis (Prentice-Hall, Englewood Cliffs, 1987).
[66] O. C. Zienkiewicz and R. L. Taylor, The finite element method: The basis, 5th edition, Vol. 1 (Butterworth and Heinemann, 2000).
[67] D. Harmon, E. Vouga, B. Smith, R. Tamstorf, and E. Grinspun, Communications of the ACM **55**(4), 102–109 (2012).
[68] M. Hauth, Visual Simulation of Deformable Models, PhD thesis, University of Tübingen, 2004.
[69] A. Lew, M. Marsden, M. Oritz, and M. West, Int. J. Numer. Meth. Engng **60**, 153–212 (2004).
[70] M. Bronstein, M. M. Bronstein, and R. Kimmel, PNAS **103**, 1168–1172 (2006).
[71] A. M. Bronstein, M. M. Bronstein, R. Kimmel, M. Mahmoudi, and G. Sapiro, IJCV **89**, 266–286 (2010).
[72] J. Sun, M. Ovsjanikov, and L. Guibas, Computer Graphics Forum **28**(5), 1383–1392 (2009).
[73] B. Dacorogna, Direct methods in the calculus of variations (Springer-Verlag, New York, 1989).
[74] N. Litke, M. Droske, M. Rumpf, and P. Schröder, An image processing approach to surface matching, in: Third Eurographics Symposium on Geometry Processing, edited by M. Desbrun and H. Pottmann (2005), pp. 207–216.
[75] J. A. Sethian, Level Set Methods and Fast Marching Methods (Cambridge University Press, 1999).
[76] J. Iglesias, B. Berkels, M. Rumpf, and O. Scherzer, A thin shell approach to the registration of implicit surfaces, in: Vision, Modeling and Visualization, edited by M. Bronstein, J. Favre, and K. Hormann (Eurographics Association, 2013), pp. 89–96.
[77] J. Iglesias, B. Berkels, M. Rumpf, and O. Scherzer, A shell model for level set matching, in preparation.
[78] P. G. Ciarlet, Mathematical Elasticity, Volume I: Three-dimensional elasticity, Studies in Mathematics and its Applications, Vol. 20 (Elsevier, 1988).
[79] M. Ovsjanikov, M. Ben-Chen, J. Solomon, A. Butscher, and L. Guibas, ACM Transactions on Graphics **31**(4), 30:1–30:11 (2012).

[80] O. Azencot, M. Ben-Chen, F. Chazal, and M. Ovsjanikov, Computer Graphics Forum **32**(5), 73–82 (2013).
[81] O. Azencot, S. Weißmann, M. Ovsjanikov, M. Wardetzky, and M. Ben-Chen, Computer Graphics Forum **33**(5) (2014), to appear.
[82] Z. Karni and C. Gotsman, Spectral compression of mesh geometry, in: Proceedings of ACM SIGGRAPH, (ACM Press, 2000), pp. 279–286.
[83] M. Ben-Chen and C. Gotsman, ACM Transactions on Graphics **24**(1), 60–80 (2005).
[84] P. Mullen, Y. Tong, P. Alliez, and M. Desbrun, Computer Graphics Forum **27**(5), 1487–1494 (2008).
[85] M. Reuter, F. E. Wolter, and N. Peinecke, Computer-Aided Design **38**(4), 342–366 (2006).
[86] M. Reuter, International Journal of Computer Vision **89**(2), 287–308 (2010).
[87] K. Hildebrandt, C. Schulz, C. V. Tycowicz, and K. Polthier, ACM Transactions on Graphics **30**(5), 119:1–119:11 (2011).
[88] C. von Tycowicz, C. Schulz, H. P. Seidel, and K. Hildebrandt, ACM Transactions on Graphics **32**(6), 213:1–213:10 (2013).
[89] E. Klassen, A. Srivastava, W. Mio, and S. H. Joshi, IEEE Transactions on Pattern Analysis and Machine Intelligence **26**(3), 372–383 (2004).
[90] M. Kilian, N. J. Mitra, and H. Pottmann, ACM Transactions on Graphics **26**(3), 64:1–64:8 (2007).
[91] M. F. Beg, M. Miller, A. Trouvé, and L. Younes, Computational anatomy: Computing metrics on anatomical shapes, in: Proceedings of 2002 IEEE ISBI, (2002), pp. 341–344.
[92] 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, (2003), pp. 635–638.
[93] P. Fletcher, C. Lu, S. Pizer, and S. Joshi, Medical Imaging, IEEE Transactions on **23**(8), 995–1005 (2004).
[94] M. Fuchs and O. Scherzer, International Journal of Computer Vision **79**(2), 119–135 (2008).
[95] P. W. Michor and D. Mumford, J. Eur. Math. Soc. **8**, 1–48 (2006).
[96] A. Srivastava, E. Klassen, S. H. Joshi, and I. H. Jermyn, Pattern Analysis and Machine Intelligence, IEEE Transactions on **33**(7), 1415–1428 (2011).
[97] G. Sundaramoorthi, A. Mennucci, S. Soatto, and A. Yezzi, SIAM Journal on Imaging Sciences **4**(1), 109–145 (2011).
[98] G. Sundaramoorthi, A. Yezzi, and A. Mennucci, International Journal of Computer Vision. **73**(3), 345–366 (2007).
[99] B. Heeren, M. Rumpf, M. Wardetzky, and B. Wirth, Computer Graphics Forum **31**(5), 1755–1764 (2012).
[100] B. Heeren, M. Rumpf, P. Schröder, M. Wardetzky, and B. Wirth, Computer Graphics Forum **33**(5) (2014), to appear.
[101] L. Rayleigh, Theory of Sound: V. 2, Band 2 (Courier Dover Publications, 1900).
[102] M. Rumpf and B. Wirth, SIAM Journal on Imaging Sciences (2014), in press.
[103] M. Rumpf and B. Wirth, IMA Journal of Numerical Analysis (2014), to appear.
[104] J. Ehlers, F. A. E. Pirani, and A. Schild, The geometry of free fall and light propagation, in: General relativity (papers in honour of J. L. Synge), (Clarendon Press, Oxford, 1972), pp. 63–84.
[105] A. Kheyfets, W. A. Miller, and G. A. Newton, Internat. J. Theoret. Phys. **39**(12), 2891–2898 (2000).