

F1 Geometry processing

Konrad Polthier, Alexander Bobenko, Klaus Hildebrandt, Ralf Kornhuber,
Christoph von Tycowicz, Harry Yserentant and Günter M. Ziegler

The field of geometry processing concerns the representation, analysis, manipulation, and optimization of geometric data. It has made rapid progress motivated by, enabling, and improving the technological possibilities for the creation of digital models from real-world objects. For example, laser scanners sample millions of points from the surface of physical objects with high accuracy and software tools produce complex digital shapes from the sampled data. This development has a strong impact on the structure of shape processing in industry. As a consequence, software systems must be adjusted to follow this trend. For example, CAD systems, which traditionally use spline representation of surfaces, need to be able to process and optimize highly resolved polygonal meshes. This creates a demand for differential geometric concepts for polygonal meshes and stable numerical, geometric, and topological algorithms.

F1—1 Deformation-based shape editing

In recent years, a special focus in geometry processing has been on schemes for deformation-based surface editing. In such a *deformation-based editing* system, see [9, 10] and references therein, a user can select parts of a geometry as handles and translate and rotate them in space. The system automatically deforms the shape so that the handles interpolate or approximate the specified positions. To provide intuitive usability, the computed deformations must be physically meaningful to match the user's intuition and experience on how shapes deform. This is achieved by computing static equilibrium states of the elastic object subject to constraints or external forces that represent the user's input. A major advantage of deformation-based editing over traditional modeling techniques, like NURBS or subdivision surfaces, is that many complex editing tasks can be described by few constraints. For example, all shapes shown in Figure 1 are created by applying one rigid transformation to three handles (the head and the two hands). This allows for efficient and simple click-and-drag user interfaces.

A challenging problem is that on the one hand to compute a deformation a non-linear optimization problem has to be solved and on the other hand a shape editing system must provide interactive response times. Hildebrandt et al. [23] developed a scheme for deformation-based editing of surface meshes based on model reduction. The scheme constructs a low-dimensional approximation of the optimization problem underlying the editing framework and thereby achieves a runtime that depends only on the complexity of the low-dimensional system. Motivated by the observation that a typical modeling session requires only a fraction of the full

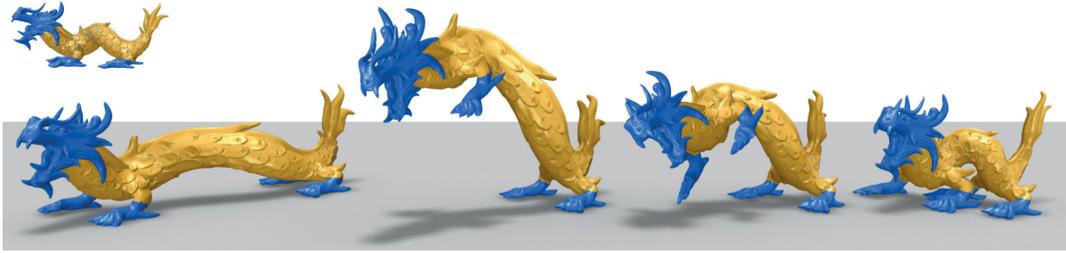


Figure 1. *Deformation-based modeling of a dragon model. Handles (blue areas) can be translated and rotated in space to define a deformation.*

shape space of a detailed mesh, they chose to apply dimension reduction to the problem. Second and third derivatives of the potential energy are used to construct a low-dimensional shape space that forms the feasible set for the optimization. For the fast approximation of the energy and its derivatives, they propose a scheme based on a second reduced shape space for a simplified mesh. By construction, the two reduced shape spaces are isomorphic and the isomorphism can be used to pull the energy from the shape space of the simplified mesh to the shape space of the full mesh. To solve the reduced optimization problem, a quasi-Newton method is used. To improve the performance, the inverse Hessian at the rest state of the energy is computed during the preprocess and used as a preconditioner for the system. Results are shown in Figures 1 and 2.

The modal reduction approach provides interactive response times, albeit at the expense of an elaborate preprocess. Recently, von Tycowicz et al. [45] proposed efficient reduction techniques for the approximation of reduced forces and for the construction of reduced shape spaces of deformable objects that accelerate the construction of a reduced dynamical system, increase the accuracy of the approximation, and simplify the implementation of model reduction. Based on the techniques, von Tycowicz et al. extend the interactive deformation-based editing scheme in [23] to elastic solids with arbitrary, nonlinear materials.

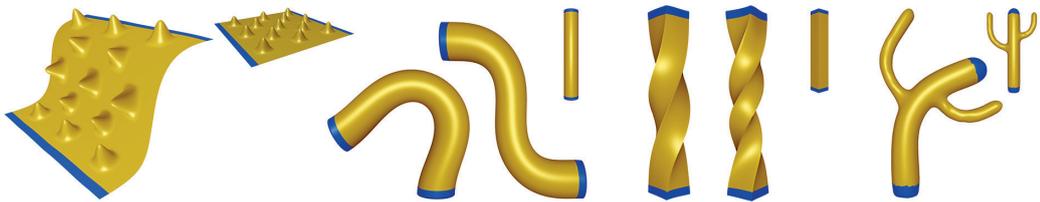


Figure 2. *Results of our geometric modeling technique are shown on the test suite of models and poses introduced in [10]. Two even larger deformations have been added.*

F1—2 Mesh fairing and smoothing

For meshes appearing in real world applications, noise is an omnipresent artefact that arises due to resolution problems in mesh acquisition processes. For example, meshes extracted from image data or supplied by laser scanning devices often carry high-frequency noise in the position of the vertices. This imposes a strong need for smoothing methods. Hildebrandt and Polthier [18] have developed a fairing method that allows to prescribe a bound on the maximum deviation of every vertex of a polyhedral surface from its initial position. The scheme is modeled as a constrained non-linear optimization problem, where a discrete fairness energy (e.g., a discrete Willmore energy) is minimized while inequality constraints ensure that the maximum deviation of the vertices is bounded. The optimization problem is solved by an active set Newton method with gradient projection.

An important application of surface smoothing is the removal of noise from 3D laser scan data. Though a laser scanner can capture the geometry of an object with high precision [32], the resulting data still contains noise. Surface smoothing methods are applied, in a post process, after a surface has been created from a number of range images. A benefit of the constraint-based fairing scheme over alternative approaches is that it can preserve the measuring accuracy of the data while smoothing out the noise. A second application of the scheme is the removal of aliasing and terracing artifacts from isosurfaces, which appear when a surface is extracted from volumetric data. It is assured that the surface remains within the domain consisting of the voxels that contain the initial surface and their 1-neighbors. In addition, the scheme was recently applied by Váša and Rus [43] for removing artifacts induced by quantization of the vertex positions, which is used for mesh compression. The fairing method offers the benefit that the vertices are kept within the cubical cells specified by the quantization.

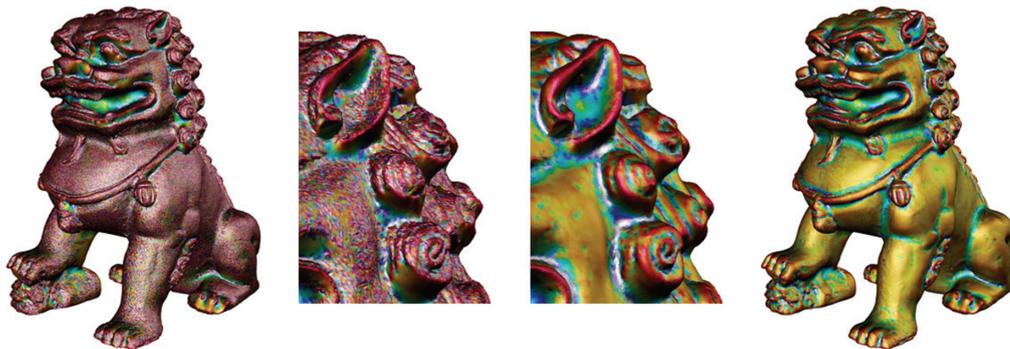


Figure 3. Left: A noisy scan of a Chinese lion with a height of 10 cm. Right: Every point of the smoothed output of our method [18] remains within a 0.1 mm distance to its initial position. The surfaces are colored by mean curvature.

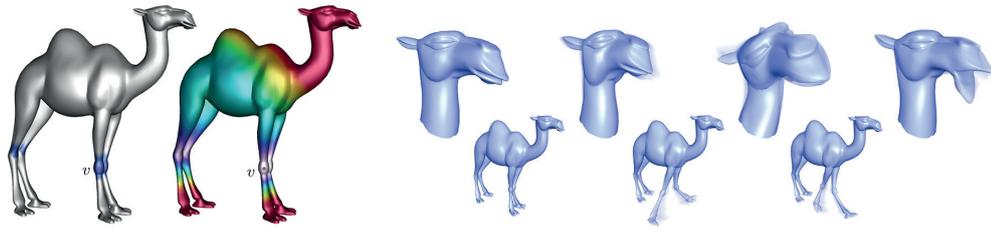


Figure 4. Results of a similarity measure that is derived from one of our shape signatures are shown on the left. Distance to the vertex v (pink dot) in binary as well as continuous coloring. The basis of the signature are vibration modes of elastic shells. Examples of vibration modes are shown on the right.

F1—3 Modal shape analysis

In recent years, substantial progress in *shape analysis* has been achieved through methods that use the spectra and eigenfunctions of discrete Laplace–Beltrami operators. Hildebrandt et al. [22, 24] have studied spectra and eigenfunctions of discrete differential operators that can serve as an alternative to discrete Laplacians for applications in shape analysis. They construct such operators as the Hessians of surface energies, which operate on a function space on the surface, or of deformation energies, which operate on a shape space of surfaces. In particular, they have designed a quadratic energy whose Hessian equals the Laplace operator if the surface is a part of the Euclidean plane. Otherwise, the Hessian eigenfunctions are sensitive to the extrinsic curvature, e.g., sharp bends, on curved surfaces. Furthermore, they considered eigenvibrations induced by deformation energies and derived a closed form representation for the Hessian (at the rest state of the energy) for a general class of such deformation energies. Based on these spectra and eigenmodes, they derive two shape signatures: one that can be used to measure the similarity of points on a surface, and another that can be used to identify features of surfaces. A conceptual difference of this similarity measure and most others is that it not only uses a local neighborhood to measure similarity, but also it uses global information encoded in the spectrum and the eigenfunctions of an adequate differential operator. This is illustrated in Figure 4, which shows an example in which our signature identifies regions of a surface as similar regions (the knees of the hind legs and the knees of the front legs of the camel) though the local geometry of the regions is different.

F1—4 Controlling dynamic shapes

Creating motions of objects or characters that are physically plausible and follow an animator’s intent is a key task in computer animation. Traditionally, the motions of objects or characters are generated from keyframes that specify values for all of the object’s or character’s degrees of freedom at a sparse set of points in time. Then, a continuous motion is obtained by fitting splines through the keyframes. This technique is attractive since it offers an adequate amount of control over the motion at a low computational cost. One drawback for this technique is that it offers little help to an animator who wants to create physically plausible motions. Physical simulation can produce realistic motions, but it is a delicate task to explicitly determine forces and physical quantities that produce a motion that matches an animator’s

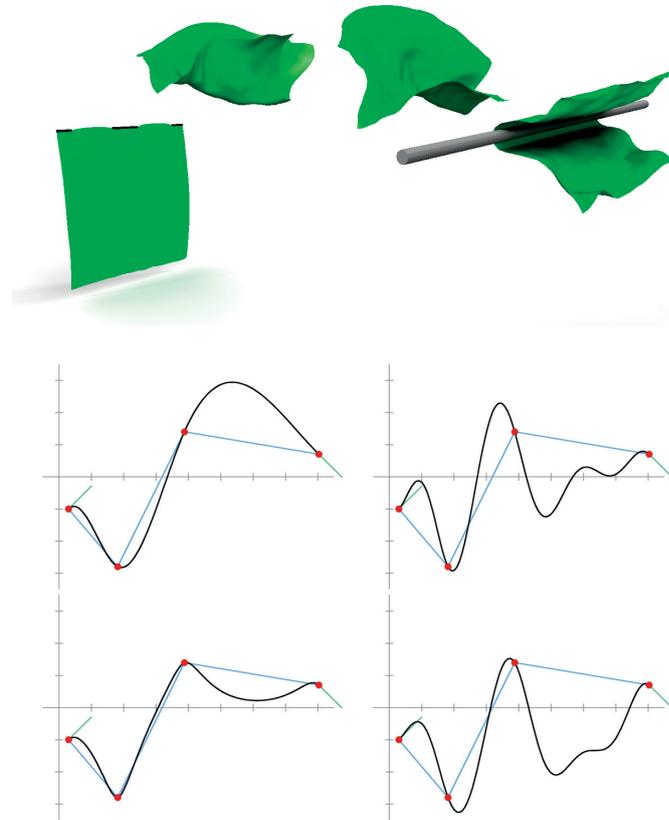


Figure 5. Top: Snapshots of a cloth animation that exhibits physical behavior but is controlled by keyframes are shown. The animation is created with our scheme for interactive spacetime control of deformable objects, see [24]. The scheme uses the concept of wiggly splines (bottom).

intentions. This is aggravated by the fact that physical simulations are integrated forward in time, which means that small changes at some point in time can have a large impact on the state of the system at a later time. Control over a simulation can be achieved by computing optimal physical trajectories that are solutions of a variational spacetime problem [48]. Such techniques calculate acting forces that minimize an objective functional while guaranteeing that the resulting motion satisfies prescribed spacetime constraints, e.g. interpolates a set of keyframes. Resulting forces are optimally distributed over the whole animation and show effects like squash-and-stretch, timing, or anticipation that are desired in animation. However, a major drawback of this approach is that a complex optimization problem must be solved to compute a motion and animators are reluctant to use any technique which slows down an animation system deviating from interactive speeds [30].

Hildebrandt et al. [24] have developed a technique for generating motions of deformable objects that can be controlled by spacetime constraints like keyframes, velocities, and forces.

The main feature of our scheme is that (after a preprocess) it offers interactive response times for creating a motion, adjusting physical parameters, or editing control parameters. This is achieved by a combination of model reduction, a multipoint linearization, modal coordinates, and a fast and robust algorithm for computing the so-called wiggly splines. Examples are shown in Figure 5.

F1—5 3D mesh compression

Compression of digital geometry models is the answer to an industrial demand: Ever-finer geometric detail, requiring millions of vertices, is part of the everyday agenda in the movie industry, the computer aided design (CAD) industry, and in server-sided rendering applications. Over the last years, many exciting ideas and new theoretical insights have been devoted to finding ways of reducing the amount of storage such models absorb. Some of those ideas have become industrial standards, like the compression methods built into MPEG-4 and Java3D. Different requirements gave rise to differing solutions with varying trade-offs between efficiency of representation and accuracy of detail – there are lossless and lossy approaches, there are wavelet and spectral decomposition methods, there are progressive as well as single-resolution techniques. But often, such as for detailed mechanical components in CAD systems, lossy storage is prohibitive, and this is where lossless coders enter. *Lossless* stands for the ability to encode the floating point positions of the mesh at highest accuracy; in practice, positions are often quantized to 10–14 bits per coordinate, a concession which has turned out to be tolerable in applications. Unlike other types of multimedia, e.g., sound and video, curved surfaces do not admit straightforward application of signal processing techniques from the Euclidean setting like the *fast Fourier transform*. However, many of these techniques can be generalized to surfaces with arbitrary topology based on the notion of semiregular meshes (also referred to as multiresolution meshes). These meshes result from successive refinement of a coarse, carefully laid out base mesh and are for example inherent to multigrid methods for solving differential equations or level-of-detail visualizations in virtual environments. Applying the refinement locally allows to increase the mesh resolution only where it is needed, however, at the expense of a non-trivial hierarchical structure. We have developed a lossless connectivity compression (see [29, 44]) that is adapted to the special characteristics of such adaptive multiresolution meshes. Using information theoretic strategies such as context-based arithmetic coding, we take advantage of structural regularities that are typically present in real-world data. Additionally, we present extensions that exploit correlations of the refinement structure in sequences of time-dependent meshes (see Figure 6). The scheme works seamlessly with wavelet-based coding strategies for which we devised improved context modeling exploiting intraband and composite statistical dependencies. This has been combined with adaptive lossy trajectory storage for adjoint gradient computation in PDE-constrained optimal control problems [16, 17, 47]. Trajectory compression was successfully applied to optimal control of cardiac defibrillation [15].

Unfortunately, in many applications 3D meshes do not possess such a hierarchical structure and therefore no assumptions can be made about its complexity, regularity or uniformity. For such irregular meshes single-rate techniques have been proven to be very efficient. In particular, our FreeLence scheme [28] belongs to this category. It uses free valences and exploits geometric information for connectivity encoding. Furthermore, FreeLence takes advantage of

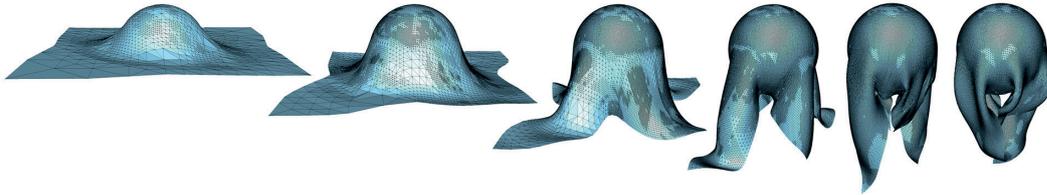


Figure 6. *The coherence of consecutive frames in time-varying sequences, as in this cloth simulation, is exploited to further improve the performance of the compression scheme.*

an improved linear prediction scheme for geometry compression of 3D meshes. Together, these approaches yield a significant entropy reduction for mesh encoding with an average of 20–30% over leading single-rate region-growing coders, both for connectivity and geometry.

F1—6 Discrete surface parametrization

For smooth surfaces, a number of special parametrizations which are adapted to the surface geometry are well known from classical differential geometry, like conformal (angle preserving) parametrizations and parametrizations by curvature lines, by asymptotic lines, by conjugate nets, etc. A particular application often demands a certain type of parametrization due to its special geometric properties. For example, conformal parametrizations are desirable for texture mapping, because in the small they scale but do not distort the texture image. Conformality is also called for when a surface is to be remeshed with nearly regular triangles or quadrilaterals. This raises very natural questions regarding discrete surfaces (meshes), like: “What does it mean for a discrete surface to be conformally parametrized?” The key challenge is to find proper discrete versions of differential geometric notions and to develop the corresponding theory. The main goal is to develop a theory of discrete surface parametrizations for arbitrary surfaces and to profitably apply it to problems arising in practice.

Discrete conformal parametrizations via circle patterns. For polyhedral surfaces, i.e., surfaces glued from planar polygons, there are various definitions of conformal maps. One definition deals with polygons inscribed in circles, and is formulated in terms of these circles. A conformal map then is a pair of circle patterns with equal intersection angles. A generalization of this definition (the angles are preserved as good as possible) was used to create conformal maps of triangulated surfaces in [31]. See also the Showcase 19 about the MATHEON bear where this method has been applied.

Discretely conformally equivalent meshes. A new conformal mesh flattening algorithm was suggested in [42] and further developed in [4], see Figure 7. It is based on a strikingly simple definition for discrete conformal equivalence: Two triangle meshes with the same combinatorics are considered discretely conformally equivalent if scale factors can be associated to the vertices such that the length of an edge in the second mesh is obtained by multiplying the length of the corresponding edge in the first mesh with the geometric mean of the scale factors associated to its two vertices. This definition discretizes in a straightforward manner

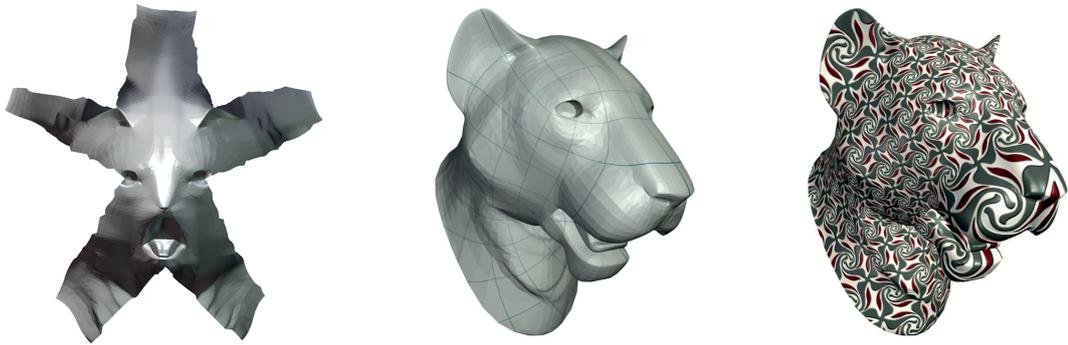


Figure 7. A discrete conformal map to the plane. Domain of parametrization (left), lines of constant parameter values (middle), seamless repeating pattern on the target geometry (right)

the concept of conformal equivalence for Riemannian metrics on a smooth manifold. Most importantly for the applications, we found a variational principle which reduces the conformal flattening problem (to find for a given surface mesh a conformally equivalent flat mesh) to an unconstrained convex optimization problem. The target function is a convex function of the (logarithmic) scale factors at vertices, whose value, gradient and Hessian can be computed efficiently. A useful feature of our method is the possibility to produce conformal parametrizations which are isometric on the boundary. This means the mesh can be flattened while the boundary edges retain their original lengths. This is desirable because we could show that among all conformal flattenings of a surface with boundary, the one with least distortion is the one that is isometric on the boundary.

Discrete quasiisothermic parametrizations. A related problem is to find a parametrization as close as possible to a conformal curvature line (isothermic) parametrization. The method suggested in [40] is based on construction of S-isothermic parametrizations for triangulated surfaces, see Figure 8. These are planar quad-meshes with touching incircles. For surfaces that do not admit discrete isothermic coordinates this method generates so called quasiisothermic parametrizations. Technically this method is an application of the conformal parametrization scheme developed in this project with special boundary conditions deduced from the principal curvature data of the surface. The parametrization along isothermic coordinates is useful for the creation of visually pleasing meshes for architectural building hulls. One of its main features is the planarity of the facets. Also the induced circle packing on the surface can be used by architects to create patterns.

F1—6.1 Low distortion parametrizations

The efficiency of multilevel methods strongly depends on the quality of the underlying hierarchy of grids. While in the solution of planar partial differential equations such a hierarchy usually comes with an adaptive refinement process, manifold meshes in geometry processing are often given as a collection of fine-grid triangles. At this stage, surface parametrizations play a crucial role as a preprocessing step for generating nested multilevel hierarchies of grids. Moreover, surface parametrization is an ongoing research topic with a wealth of other applications

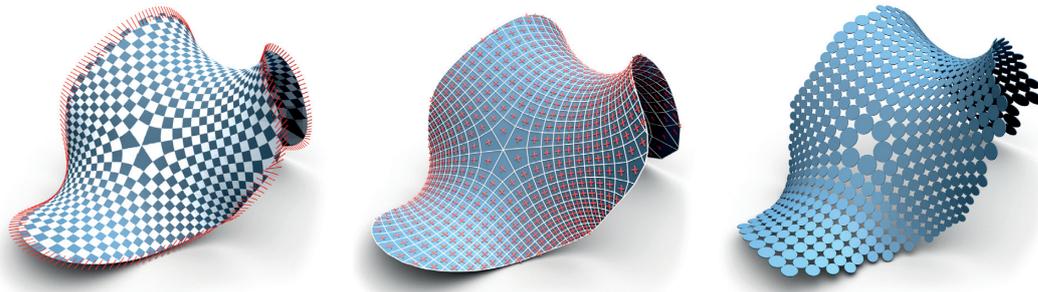


Figure 8. Discrete quasiisothermic parametrization of a piecewise flat triangulated surface. A boundary value problem is solved to create the parametrization (left), a new mesh is created from this map, edges align with the principle curvature directions of the surface (middle), the quadrilaterals possess touching incircles (right).

of their own ranging from texture mapping to extension of image processing algorithms, from remeshing to the automatic construction of hierarchical subdivision surfaces. All applications using natural coordinates will benefit from the added structure of a global parametrization.

We focus on generating a nested quadrilateral multilevel hierarchy of given triangle meshes. This hierarchy is based on the QuadCover algorithm [26], which automatically computes a quadrilateral surface mesh. QuadCover uses a curvature aligned parametrization yielding little length and area distortion of the resulting parametrization.

For controlling the alignment of quadrilaterals, a guidance frame field (e.g., derived from principal curvature directions) can be used. In a first step, the curl part of the guidance field is removed with the discrete Hodge-Helmholtz decomposition of [37], making the guidance field locally integrable. Integrating the resulting field leads to a local parametrization in the vicinity of each point, but globally, the parameter lines will not necessarily close up. Continuity is then enforced in a second step by computing a base of the first homology group and adapt the frame field for each of these base path to fulfill the closing condition.

Finally, we simplified the notion of frame fields by describing them as vector fields on a branched covering. It allows to apply the methods from classical vector field analysis to frame fields. Frame field singularities (with an index of multiples of $1/4$) appear as branch points of the covering.

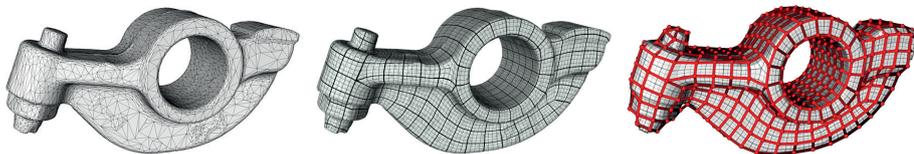


Figure 9. Automatic QuadCover parametrization (middle) from a triangle mesh (left) and a generated multigrid structure (right)

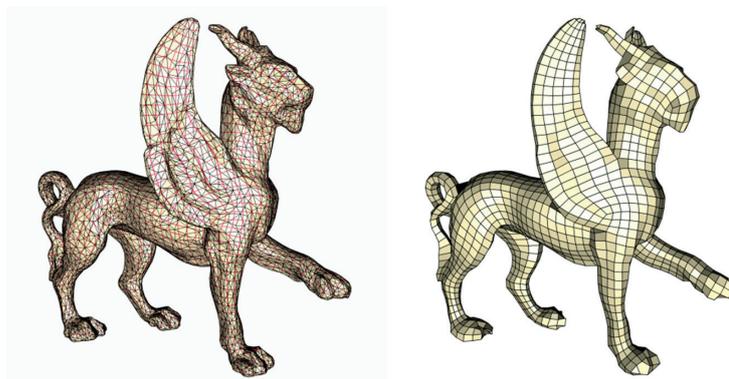


Figure 10. *Parametrization of feline model (left) and resulting quad mesh (right)*

Special care has to be taken for singularities of positive integral index. They do not resolve to branch points in the covering surface. We analysed the theory behind and gave an extension to QuadCover which allows the placement of these singularities [35].

The parameter lines divide the surface into quads which are used as a base for an adaptive nested quad hierarchy. We applied this method to generate hierarchical subdivision surfaces from irregular input meshes.

Additionally, we extended the QuadCover algorithm for stripe parametrizations of tubular objects [27]. It allows to map a regular stripe pattern globally onto tube-like surfaces such as vessels, neurons, trees, etc. Stripe parametrizations are quite useful for visualization of medical data.

F1—7 Convergence of discrete differential operators

Differential geometry studies the geometry of curved manifolds. Traditionally, the manifolds are assumed to be differentiable and techniques from calculus are used. Discrete differential geometry develops discrete notions and concepts that describe geometric properties of discrete manifolds in analogy to the smooth theory. The results in this field are heavily used for geometric computations, which are performed on discrete manifolds since computers can only process finite sets of numbers.

An important aspect of this theory is the construction of discrete differential operators and discrete curvatures on polyhedral surfaces (or polygonal meshes) and the study of their convergence properties.

Discrete Laplace–Beltrami operators. Discrete Laplace–Beltrami operators are basic objects in discrete differential geometry [7], discrete complex analysis [13, 34], and numerics of geometric partial differential equations [14]. In addition, different applications in fields like computer graphics [12, 33], geometry and image processing [9, 39], computational biology [8], and neuroscience [1, 41] use discretizations of the Laplace–Beltrami operator.

Among the different discretizations of the Laplace–Beltrami operator on polyhedral surfaces, the cotan Laplacian [36] introduced by Pinkall and Polthier, is probably the most promi-

ment. Wardetzky et al. [46] analyze structural properties of discrete Laplace–Beltrami operators. Building on the continuous setting, they propose a set of desirable properties for discrete Laplace–Beltrami operators and prove a theoretical limitation: discrete Laplacians cannot satisfy all the properties. For example, the cotan Laplacian satisfies all but one of the properties, namely the maximum principle. In addition to the analysis, Wardetzky et al. introduce a construction of discrete Laplace–Beltrami operators that uses the outer differential of discrete 1-forms. A discrete Laplace–Beltrami operator is obtained by specifying an L^2 -product on the space of discrete 1-forms. Using the concept of an intrinsic Delaunay triangulation of a polyhedral surface, Bobenko and Springborn [6] propose a modified cotan Laplacian that has non-negative weights. This implies that the discrete operator satisfies a maximum principle, which, in general, is not satisfied by the cotan Laplacian. For an example of a cotan-discrete minimal surface that does not satisfy the maximum principle, we refer to [38].

Hildebrandt et al. [21] established convergence results for a wide class of discrete differential geometric properties of polyhedral surfaces, such as convergence of geodesics, convergence of the surface area, weak convergence of the mean curvature, weak convergence of the Laplace–Beltrami operators, and convergence of solutions to the Dirichlet problem of Poisson’s equation. In particular, *convergence of the prominent cotan-formula* was shown, proving consistency of this finite-element approach with the smooth theory. Beyond these results, an important question is whether one can construct a consistent discretization of the strong form of the Laplace–Beltrami operator, i.e., a discretization that converges pointwise. Based on the cotan weights, various constructions of discrete Laplacians have been proposed. However, pointwise convergence results for these operators could only be established for special types of meshes (e.g., meshes with certain valences) and counterexamples to consistency have been reported [21, 49]. Hildebrandt and Polthier [20] introduced a discretization of the strong Laplace–Beltrami operator based on the cotan-weights and prove its consistency.

Discrete Willmore energy. The Willmore energy of a smooth surface M in \mathbb{R}^3 is the nonlinear geometric functional

$$W(M) = \int_M H^2 dvol.$$

$W(M)$ agrees, modulo multiples of the total Gauß curvature $\int_M K dvol$, with the functionals

$$\int_M (\kappa_1^2 + \kappa_2^2) dvol \quad \text{and} \quad \int_M (\kappa_1 - \kappa_2)^2 dvol. \quad (1)$$

The Gauß–Bonnet theorem implies that the total Gauß curvature is constant under variations of a surface that keep the boundary and tangent planes at the boundary fixed. Hence, under such boundary constraints a minimizer of the Willmore energy is also a minimizer of the other two functionals. This means it has the least curvature (as a minimizer of the first functional in (1)) and the least difference in the principal curvatures (second functional). In addition, the second functional in (1) has the remarkable property that it is invariant under Möbius transformations of \mathbb{R}^3 , see [2].

Boundary value problems for the Willmore energy are of fourth order, which makes discretizing the Willmore energy and the associated flow on polyhedral surfaces a difficult task. Based on a discretization of the mean curvature vector, Hsu, Kusner, and Sullivan [25] introduced a discrete Willmore energy for polyhedral surfaces and used Brakke’s Surface Evolver [11] to compute minimizers with different genus. Bobenko [3] proposed a discrete Willmore energy

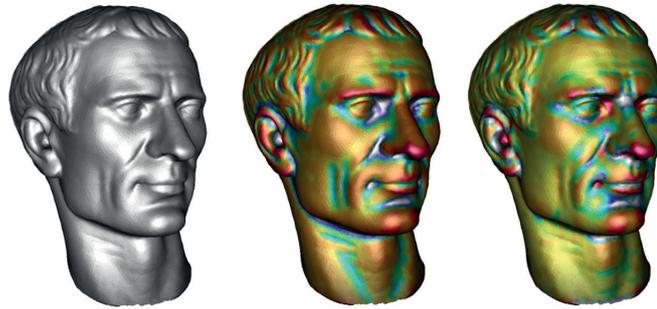


Figure 11. Mean curvature (middle) and Gaussian curvature (right) computed using a generalized shape operator on a 3d-scanned model. Color coding from white (negative) to red (positive).

for polyhedral surfaces that preserves the Möbius symmetry of the continuous energy. In [5] the flow of this discrete energy is studied.

On smooth surfaces, the Willmore energy is linked to the Laplace–Beltrami operator through the mean curvature vector field \mathbf{H} , which is the product of the mean curvature H and the surface normal field. The mean curvature vector field equals the Laplace–Beltrami operator of the embedding of the surface, thus the Willmore energy equals the squared L^2 -norm of the mean curvature vector. Since the embedding of a polyhedral surface is continuous and piecewise linear (hence in the domain of the discrete Laplace–Beltrami operators), the construction of the discrete strong Laplace–Beltrami operators in [20] extends to a construction of discrete mean curvature vectors and discrete Willmore energies. Pointwise approximation of the mean curvature vector field of a surface and consistency of the discrete Willmore energies were proved.

Convergence of discrete curvatures. Curvature is a central concept in the study of geometric properties of surfaces in \mathbb{R}^3 and appears in many interesting geometric and physical problems. Examples are the study and construction of surface with constant mean curvature and the analysis and integration of curvature flows. The estimation of curvatures of a smooth surface from an approximating discrete surface is important for the numerical treatment of such problems and for various applications in engineering and computer graphics.

In classical differential geometry, the curvatures of a smooth surface M in \mathbb{R}^3 are represented by the shape operator S , a tensor field on the tangent bundle of M . Since the definition of S involves second derivatives of the embedding of the surface, it does not apply to polyhedral surfaces. A polyhedral surface has planar triangles and its *curvatures* are concentrated at the edges and vertices. Hence, roughly speaking, they cannot be described by functions but by distributions. Hildebrandt and Polthier [19] implemented this idea by introducing generalized shape operators that can be rigorously defined for smooth and polyhedral surfaces. The generalized shape operators are functionals on an appropriate Sobolev space of weakly differentiable vector fields. They showed that this description of curvature of polyhedral surfaces can be used for the pointwise approximation of the classical shape operator of a smooth surface. These are the first pointwise approximation results for the shape operator of a smooth surface from polyhedral surfaces in this generality.

References

- [1] M. Belkin and P. Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Comput.*, 15(6):1373–1396, 2003. URL: <http://dx.doi.org/10.1162/089976603321780317>.
- [2] W. Blaschke. *Vorlesungen über Differentialgeometrie III: Differentialgeometrie der Kreise und Kugeln*, volume 29 of *Grundlehren der Mathematischen Wissenschaften*. Springer, 1929.
- [3] A. I. Bobenko. A conformal energy for simplicial surfaces. *Combinatorial and Computational Geometry*, 52:133–143, 2005.
- [4] A. I. Bobenko, U. Pinkall, and B. Springborn. Discrete conformal maps and ideal hyperbolic polyhedra. Preprint arXiv:1005.2698v1, 2010.
- [5] A. I. Bobenko and P. Schröder. Discrete Willmore flow. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Geometry Processing*, pages 101–110, 2005.
- [6] A. I. Bobenko and B. A. Springborn. A discrete Laplace–Beltrami operator for simplicial surfaces. *Discrete Comput. Geom.*, 38(4):740–756, 2007. URL: <http://dx.doi.org/10.1007/s00454-007-9006-1>.
- [7] A. I. Bobenko, J. M. Sullivan, P. Schröder, and G. M. Ziegler. *Discrete Differential Geometry*. Oberwolfach Seminars. Birkhäuser Basel, 2008.
- [8] A. Bonito, R. H. Nochetto, and M. Sebastian Pauletti. Parametric FEM for geometric biomembranes. *J. Comput. Phys.*, 229:3171–3188, 2010. URL: <http://dx.doi.org/10.1016/j.jcp.2009.12.036>.
- [9] M. Botsch, L. Kobbelt, M. Pauly, P. Alliez, and B. Levy. *Polygon Mesh Processing*. AK Peters, 2010.
- [10] M. Botsch and O. Sorkine. On linear variational surface deformation methods. *IEEE Transactions on Visualization and Computer Graphics*, 14(1):213–230, 2008.
- [11] K. Brakke. The surface evolver. *Experimental Mathematics*, 1(2):141–165, 1992.
- [12] M. Desbrun, E. Grinspun, P. Schröder, and M. Wardetzky. Discrete Differential Geometry: An applied introduction. In *ACM SIGGRAPH ASIA Courses*, 2008.
- [13] R. J. Duffin. Potential theory on a rhombic lattice. *Journal of Combinatorial Theory*, 5(3):258–272, 1968. doi : 10.1016/S0021-9800(68)80072-9.
- [14] G. Dziuk. Finite elements for the Beltrami operator on arbitrary surfaces. In *Partial Differential Equations and Calculus of Variations*, volume 1357 of *Lec. Notes Math.*, pages 142–155. Springer, 1988.
- [15] S. Götschel, C. Nagaiah, K. Kunisch, and M. Weiser. Lossy compression in optimal control of cardiac defibrillation. *J. Sci. Comput.*, to appear, 2013. doi : 10.1007/s10915-013-9785-x.
- [16] S. Götschel, C. von Tycowicz, K. Polthier, and M. Weiser. Reducing memory requirements in scientific computing and optimal control. *ZIB Report*, (13-64), 2013.
- [17] S. Götschel and M. Weiser. Lossy compression for PDE-constrained optimization: Adaptive error control. *ZIB Report*, (13-27), 2013.
- [18] K. Hildebrandt and K. Polthier. Constraint-based fairing of surface meshes. In *Siggraph/Eurographics Sympos. Geom. Processing*, pages 203–212, 2007.
- [19] K. Hildebrandt and K. Polthier. Generalized shape operators on polyhedral surfaces. *Computer Aided Geometric Design*, 28(5):321–343, 2011. doi : DOI : 10.1016/j.cagd.2011.05.001.
- [20] K. Hildebrandt and K. Polthier. On approximation of the Laplace–Beltrami operator and the Willmore energy of surfaces. *Computer Graphics Forum*, 30(5):1513–1520, 2011. URL: <http://dx.doi.org/10.1111/j.1467-8659.2011.02025.x>.
- [21] K. Hildebrandt, K. Polthier, and M. Wardetzky. On the convergence of metric and geometric properties of polyhedral surfaces. *Geometricae Dedicata*, 123:89–112, 2006.
- [22] K. Hildebrandt, C. Schulz, C. von Tycowicz, and K. Polthier. Eigenmodes of surface energies for shape analysis. In B. Murrain, S. Schaefer, and G. Xu, editors, *Advances in Geometric Modeling and Processing*, volume 6130 of *Lecture Notes in Computer Science*, pages 296–314. Springer, 2010.

- [23] K. Hildebrandt, C. Schulz, C. von Tycowicz, and K. Polthier. Interactive surface modeling using modal analysis. *ACM Trans. Graph.*, 30:119:1-119:11, October 2011. URL: <http://doi.acm.org/10.1145/2019627.2019638>.
- [24] K. Hildebrandt, C. Schulz, C. von Tycowicz, and K. Polthier. Interactive spacetime control of deformable objects. *ACM Trans. Graph.*, 31(4):71:1-71:8, July 2012. URL: <http://doi.acm.org/10.1145/2185520.2185567>.
- [25] L. Hsu, R. Kusner, and J. Sullivan. Minimizing the squared mean curvature integral for surfaces in space forms. *Experimental Mathematics*, 1(3):191-207, 1992.
- [26] F. Kälberer, M. Nieser, and K. Polthier. QuadCover - surface parametrization using branched coverings. *Computer Graphics Forum*, 26(3):375-384, 2007.
- [27] F. Kälberer, M. Nieser, and K. Polthier. Stripe parameterization of tubular surfaces. *Proceedings TopoInVis 09, Submitted*, 2009.
- [28] F. Kälberer, K. Polthier, U. Reitebuch, and M. Wardetzky. Freelence - coding with free valences. *Computer Graphics Forum*, 24(3):469-478, 2005.
- [29] F. Kälberer, K. Polthier, and C. von Tycowicz. Lossless compression of adaptive multiresolution meshes. In *Proceedings*, 2009.
- [30] M. Kass and J. Anderson. Animating oscillatory motion with overlap: wiggly splines. *ACM Trans. Graph.*, 27:28:1-28:8, 2008. URL: <http://doi.acm.org/10.1145/1360612.1360627>.
- [31] L. Kharevych, B. Springborn, and P. Schröder. Discrete conformal mappings via circle patterns. *ACM Trans. Graph.*, 25(2):412-438, 2006.
- [32] M. Levoy, K. Pulli, B. Curless, S. Rusinkiewicz, D. Koller, L. Pereira, M. Ginzton, S. Anderson, J. Davis, J. Ginsberg, J. Shade, and D. Fulk. The digital Michelangelo project: 3D scanning of large statues. In *Proceedings of ACM SIGGRAPH*, pages 131-144, 2000. URL: <http://dx.doi.org/10.1145/344779.344849>.
- [33] B. Lévy and H. Zhang. Spectral mesh processing. In *ACM SIGGRAPH ASIA Courses*, pages 1-47, 2009.
- [34] C. Mercat. Discrete Riemann surfaces and the Ising model. *Comm. Math. Phys.*, 218(1):177-216, 2001.
- [35] M. Nieser and K. Polthier. Parameterizing singularities of positive integral index. In *13th IMA International Conference on Mathematics of Surfaces XIII*, pages 265-277, 2009.
- [36] U. Pinkall and K. Polthier. Computing discrete minimal surfaces and their conjugates. *Experim. Math.*, 2:15-36, 1993.
- [37] K. Polthier and E. Preuss. Identifying vector field singularities using a discrete Hodge decomposition. In *Visualization and Mathematics III*, pages 113-134. Springer Verlag, 2003.
- [38] K. Polthier and W. Rossmann. Counterexample to the maximum principle for discrete minimal surfaces. Electronic Geometry Model No. 2000.11.040, 2001. www.eg-models.de/2000.11.040.
- [39] M. Reuter, S. Biasotti, D. Giorgi, G. Patanè, and M. Spagnuolo. Discrete Laplace-Beltrami operators for shape analysis and segmentation. *Computers & Graphics*, 33(3):381-390, 2009. URL: <http://dx.doi.org/10.1016/j.cag.2009.03.005>.
- [40] S. Sechelmann, T. Rörig, and A. I. Bobenko. Quasiisothermic mesh layout. In L. H. et al., editor, *Advances in Architectural Geometry 2012*, pages 243-258. Springer Vienna, 2013. URL: http://dx.doi.org/10.1007/978-3-7091-1251-9_20.
- [41] H. Sprekeler. On the relation of slow feature analysis and Laplacian eigenmaps. *Neural Comput.*, 23(12):3287-3302, 2011. URL: http://dx.doi.org/10.1162/NECO_a_00214.
- [42] B. Springborn, P. Schröder, and U. Pinkall. Conformal equivalence of triangle meshes. *ACM Trans. Graph.*, 27(3), 2008. URL: <http://doi.acm.org/10.1145/1360612.1360676>.
- [43] L. Váša and J. Rus. Dihedral angle mesh error: a fast perception correlated distortion measure for fixed connectivity triangle meshes. *Computer Graphics Forum*, 31(5):1715-1724, 2012.

- [44] C. von Tycowicz, F. Kälberer, and K. Polthier. Context-based coding of adaptive multiresolution meshes. *Computer Graphics Forum*, 30(8):2231–2245, 2011. URL: <http://dx.doi.org/10.1111/j.1467-8659.2011.01972.x>.
- [45] C. von Tycowicz, C. Schulz, H.-P. Seidel, and K. Hildebrandt. An efficient construction of reduced deformable objects. *ACM Trans. Graph.*, 32(6):213:1–213:10, Nov. 2013. URL: <http://doi.acm.org/10.1145/2508363.2508392>.
- [46] M. Wardetzky, S. Mathur, F. Kälberer, and E. Grinspun. Discrete Laplace operators: No free lunch. In *Siggraph/Eurographics Sympos. Geom. Processing*, pages 33–37, 2007.
- [47] M. Weiser and S. Götschel. State trajectory compression for optimal control with parabolic PDEs. *SIAM J. Sci. Comput.*, 34(1):A161–A184, 2012.
- [48] A. Witkin and M. Kass. Spacetime constraints. *Proc. of ACM SIGGRAPH*, 22:159–168, 1988. URL: <http://doi.acm.org/10.1145/378456.378507>.
- [49] G. Xu. Convergence of discrete Laplace-Beltrami operators over surfaces. *Comput. Math. Appl.*, 48:347–360, 2004.