Robust Fitting of Subdivision Surfaces for Smooth Shape Analysis

V. Estellers[1,2], F. R. Schmidt[1,3], D. Cremers[1]
[1]Technical University Munich, [2]Toshiba Research, [3]Bosch Center for Artificial Intelligence
virginia.estellers@gmail.com, frank.r.schmidt@de.bosch.com, cremers@tum.de

Abstract

Most shape analysis methods use meshes to discretize the shape and functions on it by piecewise linear functions. Fine meshes are then necessary to represent smooth shapes and compute accurate curvatures or Laplace-Beltrami eigenfunctions at large computational costs. We avoid this bottleneck by representing smooth shapes as subdivision surfaces and using the subdivision scheme to parametrize smooth surface functions with few control parameters.

We propose a model to fit a subdivision surface to input samples that, unlike previous methods, can be applied to noisy and partial scans from depth sensors. The task is formulated as an optimization problem with robust data terms and solved with a sequential quadratic program that outperforms the solvers previously used to fit subdivision surfaces to noisy data. Our experiments show that the compression of a subdivision representation does not affect the accuracy of the Laplace-Beltrami operator and allows to compute shape descriptors, geodesics, and shape matchings at a fraction of the computational cost of mesh representations.

1. Introduction

Many shape analysis tasks describe shapes as smooth manifolds and analyze them with respect to their geometry in terms of normals, curvatures, and descriptors that require accurate estimates of first- and second-order derivatives over the shape to compute tangent spaces and Laplace-Beltrami operators [17,18,27,32]. When the surface is represented as a mesh or point cloud, these differential operators can only be approximated at large memory and computational costs.

Once normals and tangent spaces have been computed, most of the high resolution information is discarded by operating only with the leading eigenfunctions of the Laplace-Beltrami operator to compute geodesics [6], shape descriptors [1,17,18,27] or shape matches [24] with tractable problem sizes. This creates a paradox, as shapes are first discretized with fine meshes or dense pointclouds to estimate high-dimensional differential operators that are mostly discarded to analyze the shape. However, we only need large meshes or dense pointclouds to discretize accurate differential operators, not for shape analysis. We avoid this paradox by representing a shape with a subdivision surface [23] that parametrizes the surface with a small set of smooth base functions. Differentiability is then intrinsic to the shape representation and does not require large bases to compute shape descriptors, geodesics, or shape matches.

Subdivision surfaces are not the surface representation obtained from range measurements or computer vision algorithms. Their impact on computer vision depends on the availability of robust techniques to fit a subdivison surface to raw measurements. For this reason, our first contribution (Section 4) is a robust method and code that fits a subdivision surface to a noisy pointcloud and estimates the geometric operators for shape analysis. Our model can also handle partial scans by reconstructing surfaces with a boundaries and is robust to outliers. When we fit a subdivision surface to triangular meshes common in shape analysis [2], our robust model improves upon least-squares techniques [10,16,21,22] sensitive to erroneous correspondences that ignore boundary information required to fit partial scans.

Our second contribution is a fast optimization algorithm (Section 6) that combines a sequential quadratic approximation with a bespoken initialization that estimates the surface topology. Initialization impacts the accuracy of the solution because the problem is not convex. As a result, [4,11,29] resort to manual or spherical [12] initializations that constraint the topology and oversmooth the surface to avoid small-scale minima and prevent self-intersections during optimization.

In our third contribution (Section 3) we demonstrate the accuracy of shape analysis with subdivision surfaces. We focus on Laplace-Beltrami operator, wave kernel signatures [1], geodesics [6] and shape matching [24]. Our experiments show state-of-the-art performance at a fraction of the memory costs of triangular meshes, see Fig. 1 .

2. Related Work: Clean Graphics, Noisy Vision

Graphics techniques fit a subdivision surface by minimizing the sum of squared distances between the subdivision surface and the vertices of a clean mesh [10, 21]. Research has focused into efficient ways of minimizing this

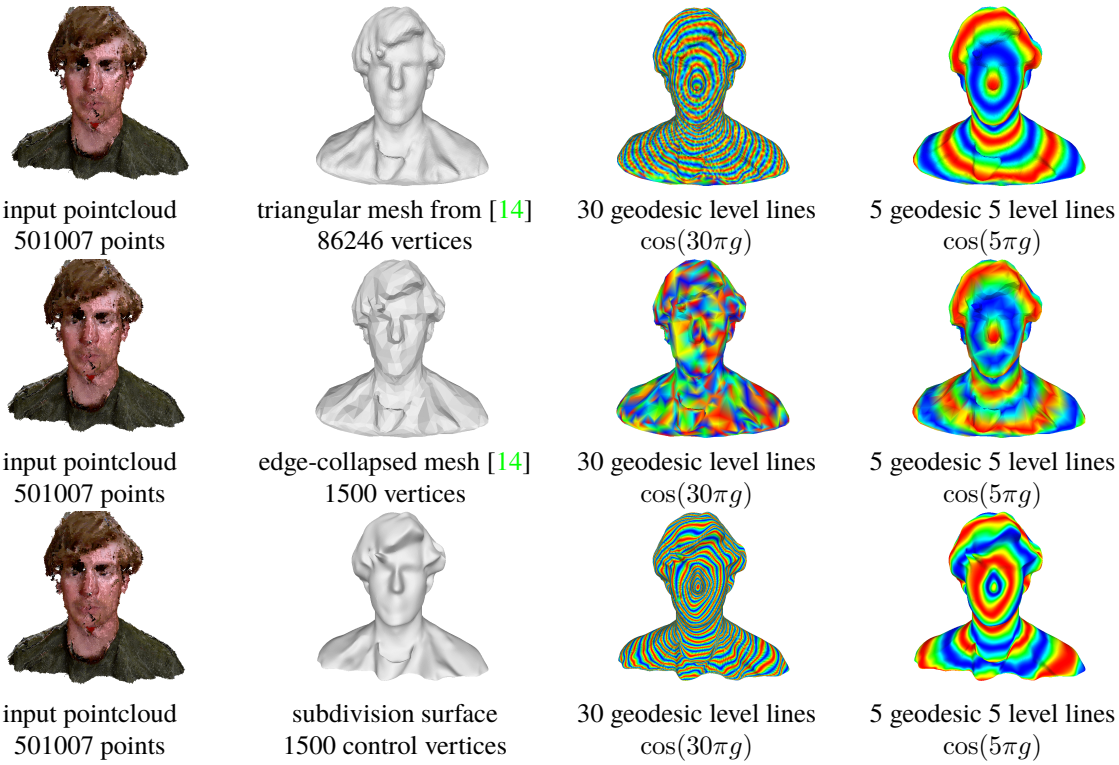

Figure 1. Teaser: Computation of geodesic g with different surface representations by the Heat method [6] . Column 1: input Kinect pointcloud. Column 2: representation of the surface with a high-resolution mesh obtained with Poisson reconstruction [14] (row 1), a low-resolution mesh obtained by quadratic edge collapse [8] of the high-resolution mesh (row 2), and with our subdivision surface (row 3). Columns 3 and 4: level lines of the geodesic visualized as $\cos(\varpi g)$. Our subdivision surface is comparable to the surface obtained by Poisson reconstruction and faithfully represent geodesics at low and high resolution with 2% of vertices. Compressing the Poisson mesh by edge collapse looses small scale details of the surface and its geodesic, as highlighted by high-frequency level lines of the geodesic.

non-convex energy by gradient-descent [33], quasi-Newton [19, 22, 30], and sequential convex techniques [16, 22]. A common approach iteratively approximates the squared distance to the surface and minimizes the resulting convex energy. Our method adapts this strategy to the robust energies that fit a surface to noisy data. We outperform these models in 4 points: 1) A robust energy reduces the effects of outliers in the input data, errors in the correspondences (between input samples and the surface point closest them), and the quadratic distance approximation. 2) We use a 2nd-order approximation to the squared distance function, instead of the 1st order of [22], to take into account curvature information and fit better the small-scale structures. 3) We penalize deviations of the tangent space to improve the appearance of sharp creases. 4) We fit surfaces with boundaries.

Computer Vision has focused on energy models to fit subdivision surfaces to noisy data [4, 11, 12, 29], but has ignored the impact of outliers and optimization. [11] fits a subvision surface to range data by solving a least-squares problem that assumes optimal correspondences. This assumption is removed in [4, 12, 29] by explicitly optimizing the correspondences. These methods assume the topology of the surface to be know and initialize it manually with a correct but coarse mesh [4, 29] or with a sphere [12] that needs to oversmooths the energy (surface) to prevent self intersections during optimization. The Levenberg-Marquardt (LM) solvers of these methods estimate jointly the vertices of the control mesh and the correspondences by increasing the size of the problem. Without frequent reinitialization, LM solvers are slow because the correspondence updates are limited to to adjacent faces of the control mesh. Our algorithm places no limits on these updates and can jump to different valleys of the energy landscape.

To our knowledge subdivision surfaces have not been used in shape analysis before. [3, 9, 13, 31] investigate the accuracy of mass and stiffness matrices on subdivision but do not consider shape descriptors, geodesics, or matching.

3. Smooth Shapes

A shape is a smooth (C^2), compact, oriented surface $S \subset \mathbb{R}^3$ with boundary ∂S. When we discretize a smooth surface by a mesh $\mathbf{M} = (V, F)$, with vertices V and facets F, the surface parametrization is only piecewise C^1 because each facet parametrizes a piece of surface independently.

Subdivision surfaces, on the other hand, use meshes to discretize the domain of the parametrization, not the surface itself, and obtain a smooth discretization without explicitly constraining the parametrization to be smooth across surface facets. We refer to [23] for an introduction to subdivision.

A subdivision surface represents a smooth surface S by a coarse control mesh $\mathbf{M} = \mathbf{M}^0$ and a subdivision scheme. The subdivision scheme transforms a mesh $\mathbf{M}^{k-1}$ into a finer mesh $\mathbf{M}^k$ and obtains a smooth surface as the limit of iterating subdivision ad infinitum, that is $S = \lim_{k\to\infty} S_{\mathbf{M}^k}$.

The topology of the surface is completely determined by the topology of the control mesh M^0, while its geometry is characterized by the location of the its vertices $V = (v_1 ... v_n)$. The limit surface S is parametrized by a linear combination of the limit base functions of the subdivision scheme $\Phi_1, \ldots, \Phi_n$ with the control vertices as follows

$$X(u) = \sum_{i=1}^{n} v_i \Phi_i(u) \;\text{ for }\; u = (\hat{u}, f) \in \mathcal{T}_{\mathbf{M}}, \tag{1}$$

where the domain of the parametrization, $\mathcal{T}_{\mathbf{M}}$, is the set of mesh faces $f \in F^0$ with local coordinates $\hat{u}$. The limit base functions $\Phi_1, \ldots, \Phi_n$ are compactly supported because the subdivision rules act locally and correspond to different spline basis [23]. The equivalence of the limit functions to spline bases provides fast evaluation techniques [28] for functions and tangent vectors defined over the surface. This makes subdivision well-suited for shape analysis.

We adopt the Catmull-Clark [5] subdivision scheme that generalizes bi-cubic uniform B-splines to surfaces of arbitrary topology because it has a simple tensor-product structure and the limit surface is C^2 everywhere except at extraordinary vertices of with valence $\neq 4$, where it is are at least C^1. The scheme is designed for control grids with quadrilateral connectivity, but our method can easily be applied to other schemes with the same smoothness.

4. Robust Fitting Model

Given the surface samples $\mathbf{P} = \{p_1 ... p_N\}$ with associated normals $\mathbf{T} = \{t_1 ... t_N\}$ and the boundary samples $\bar{\mathbf{P}} = \{\bar{p}_1, \ldots, \bar{p}_m\}$, we fit a subdivision surface $S = \{X(u) = \sum_{i=1}^{n} \Phi_i(u) v_i \mid u \in \mathcal{T}_{\mathbf{M}}\}$ to these samples by solving the constrained minimization problem:

$$\min_{S} \mathcal{E}(S) \;\text{ s.t. }\; \partial S \ni \bar{p_j} \text{ for } j \leq m, \tag{2}$$
$$\text{where }\; \mathcal{E}(S) = \text{dist}(S, \mathbf{P}) + \alpha \cdot \mathcal{T}(S, \mathbf{T}) + \beta \cdot R(S),$$

$\alpha, \beta \geq 0$ are model parameters, $\text{dist}(S, \mathbf{P})$ fits the surface S to the points $\mathbf{P}$, $\mathcal{T}(S, \mathbf{T})$ makes the tangent plane of S orthogonal to the normals $\mathbf{T}$, $R(S)$ regularizes its control mesh, and the constraints ensure that boundary points map to surface boundaries. The subdivision surface $S(V, F)$ is parametrized by its control mesh $\mathbf{M} = (V, F)$.

Point Fit Let $\mathcal{D}_S(p) = \min_{s\in S} \|s - p\|$ be the Euclidean distance between a point p and the surface S, then

$$\text{dist}(S, \mathbf{P}) := \sum_{j=1}^{N} \mathcal{D}_S(p_j) = \sum_{j=1}^{N} \min_{u_j \in \mathcal{T}_{\mathbf{M}}} \|X(u_j) - p_j\|$$

Penalizing the distance between the surface and the point samples, instead of the common squared distance, makes our model robust to outliers and geometrically meaningful that the Manhattan distance associated to a robust ℓ_1 penalty. The proposed penalty is less robust than M-estimators (e.g. biweight loss), but is easier to optimize; we believe it offers a better trade off between robustness and optimization complexity for applications with a moderate number of outliers like shape analysis. The minimization in $u_1, \ldots, u_N$ results from the parametric representation of the surface $S = X(\mathcal{T}_{\mathbf{M}}) = \{X(u) \mid u \in \mathcal{T}_{\mathbf{M}}\}$ and introduces a large number of additional variables over which the objective function is not convex. We will avoid optimizing explicitly the correspondence parameters $u_1, \ldots, u_N$ by using a second-order approximation of the distance to a surface.

Tangent Fit $\mathcal{T}$ aligns the tangent space of the surface with the direction orthogonal to the input normals with

$$\mathcal{T}(S, \mathbf{T}) = \sum_{j=1}^{N} |\langle t_j, \partial_1 \Phi(u_j) \rangle| + |\langle t_j, \partial_2 \Phi(u_j) \rangle|,$$

where $\partial_i X(u_j)$ the i-th basis of the tangent space $T_{X(u_j)}S$. $\mathcal{T}$ is by design independent of the normal orientation because estimating consistent orientations in noisy point clouds is prone to fail. The ℓ_1-norm makes this term robust to errors in the correspondence parameters $u_1 ... u_N$ and outliers $\mathbf{T}$.

Boundary Constraints For each point in the surface boundary $\bar{p_j}$, the constraint $\bar{p_j} \in \partial S$ ensures that it is mapped to a point in the surface boundary. As points on the surface boundary correspond to points $X(\bar{u}_j)$ with $\bar{u}_j \in \partial\mathcal{T}_{\mathbf{M}}$, the constraints are equivalent to

$$\bar{p_j} = X(\bar{u}_j) = \sum_{i=1}^{n} \Phi_i(\bar{u}_j) v_i \;\text{ where }\; \bar{u}_j \in \partial\mathcal{T}_{\mathbf{M}} \;\; j = 1, \ldots, m.$$

Regularization The regularizer $R(S)$ penalizes the squared distance between the vertices of the control mesh incident to the same quad. This keeps the size and shape of quads regular and avoids skewed elements unstable in finite-element (FE) computations. $R(S) = \|RV\|^2$ is a simple quadratic penalty described by a sparse matrix R that has a row for each edge $e = (i, j)$ in $\mathbf{M}^0$ with entries $1, -1$ at the i, j-th columns associated with the vertices v_i and v_j incident to e.

Color Representation When the input point cloud has color, we estimate a color field $h\colon S \to \mathbb{R}^3$ over the surface parametrized with the same limit base functions by means

of the pullback $h^*\colon \mathcal{T}_{\mathbf{M}} \to \mathbb{R}^3$. The coefficients $H = (h_1, ..., h_n)$ of the color field $h^*(u) = \sum_{i=1}^n \Phi_i(u) h_i$ solve

$$\min_{h_1,\dots,h_n} \sum_{j=1}^{N} \min_{u_j \in \mathcal{T}_{\mathbf{M}}} \|h^*(u_j) - c_j\|^2 + \alpha \|X(u_j) - p_j\| + \gamma \|RH\|^2,$$

where c_j is the colors associated with p_j and $\gamma \geq 0$ a model parameter. The term $\min_{u_j} \|c_j – h^*(u_j)\|^2$ penalizes the squared distance between c_j and the color of the surface point $X(u_j)$ closest to p_j, and $\|RH\|^2$ regularizes the color fields H by penalizing color differences between adjacent control vertices. We can estimate h^* as we reconstruct the S to exploit the joint terms $\|X(u_j) - p_j\|$ in the optimization.

5. Initialization and Topology Estimation

Initialization is important because the optimization problem is not convex, the initial control mesh determines the reachable local minima, and constrains the surface topology.

The variables of our optimization are the vertices and facets of the control mesh $\mathbf{M}$ that parametrizes the surface. To avoid solving a combinatorial problem, we fix the topology of $\mathbf{M}$, and thus the surface topology, and optimize the location of the vertices V. The topology of $\mathbf{M}$ is designed to match the topology of the surface but have a fewer vertices.

If $\mathbf{P}$ are the vertices of an input high-resolution mesh $\overline{\mathbf{M}}$, the topology of the surface and its boundary samples $\bar{\mathbf{P}}$ are encoded by $\overline{\mathbf{M}}$. We split the faces of $\overline{\mathbf{M}}$ into triangles and use quadratic edge-collapse [8] to reduce the number of vertices while preserving topology. We then transform the collapsed triangular mesh into a quad mesh by solving a perfect matching problem that pairs triangles to create quads [26]. For each edge e, let $\alpha_1(e), \dots, \alpha_4(e) \in \mathbb{S}^1$ be the angles of the quad that results from removing e and $\eta(e) \in \mathbb{S}^1$ the angle between the normals of the triangles incident to e. We favour rectangular quads and avoid skewed elements (unstable to in FE computations) by setting the cost of removing an edge to

$$c(e) = \frac{1}{4} \sum_{i=1}^{4} \left(\alpha_i(e) - \frac{\pi}{2} \right)^2 + \tan(\eta(e))^2$$

with ∞ cost for edges that connect strongly bent triangles ($2\eta(e) > \pi$) or create quads with straight or reflex angles. To support the existence of a perfect matching, we partition triangles with boundary corners or in meshes with odd number of faces and solve the perfect matching problem with [15].

When the input is a point cloud, we use an implicit representation to estimate the topology of the surface from the points and only then extract a mesh. Estimating first the topology with an implicit representation avoids handling topology changes during fitting. We use Poisson reconstruction [14] but other techniques work well. From this implicit representation, we extract a triangular mesh with marching cubes [20], determine the boundary samples $\bar{\mathbf{P}}$ from its boundary vertices, and create a compact quad mesh with the collapse and blossom techniques used for meshes.

6. Efficient Optimization

Our optimization algorithm exploits the properties of subdivision surfaces, namely, the compact support of the basis functions and the ability to analytically evaluate its geometry, with a sequential quadratic program that is more efficient than the gradient-based algorithms of [4, 11, 12]. We solve (2) by solving a sequence of convex problems that approximate the original energy $E(S)$ around the current surface estimate S^t with the constrained least-squares problem

$$v^{t+1} \leftarrow \min_{v} v^\top \left(Q^\top v - b \right) \text{ s.t. } Cv = \bar{\mathbf{P}} \qquad (3)$$

The KKT conditions of (3) result in the linear system

$$\begin{pmatrix} Q & C^\top \\ C & 0 \end{pmatrix} \begin{bmatrix} v \\ \lambda \end{bmatrix} = \begin{bmatrix} b \\ \bar{P} \end{bmatrix}$$

that we can solve efficiently with an iterative linear solver warm-started with the solution from iteration t because the matrix is sparse as a result of the compact support of the subdivision basis. We explain how to derive the approximation (3) from the surface geometry without color for simplicity.

We follow a Majorize-Minimize (MM) principle to minimize $E(S)$ by iterating two steps until convergence. The first step of iteration t finds a majorizer of the objective function $E(S|S^t) \geq E(S)$ that coincides with E at S^t. In particular, our majorizer results from applying the inequality

$$|d|^q \leq \frac{q}{2} |d_0|^{q-2} d^2 + (1 - \frac{q}{2}) |d_0|^q \quad \forall d_0 \neq 0,\ q \in [1, 2]$$

with $q = 1$ for the point and tangent fit terms, that is,

$$E(S|S^t) = \sum_{j=1}^{n} w_j^t \mathcal{D}_S(p_j)^2 + \sum_{i=1}^{2} \alpha_{ij}^t \langle t_j, \partial_i \Phi(u_j) \rangle^2 + \beta R(S).$$

The second step of iteration t drives the value of the original function downwards by minimizing this upper envelop. This defines a weighted least-squares problem where the weights w_j^t and α_{ij}^t are determined by the MM principle. To convexify the problem, we approximate $\mathcal{D}_S(\cdot)^2$ with a quadratic function that parametrizes the squared distance to a surface in terms of its geometry as follows:

$$\mathcal{D}_S^2(p) \approx (x-p)^\top \left[\frac{d \cdot \tau_1 \tau_1^\top}{d + \rho_1} + \frac{d \cdot \tau_2 \tau_2^\top}{d + \rho_2} + \nu \nu^\top \right] (x-p), \quad (4)$$

where x is a point on the surface S close to p (but not necessarily the projection of p onto S), $d := \|x - p\|$ is the distance from x to this point, and ρ_i, τ_i, ν are the principal curvature radii, the principal curvature directions and the normal to the surface at p. The approximation, adapted from [25], coincides with the second-order Taylor approximation of the squared distance within the radius of curvature of the surface at x, where it can be relaxed into a positive definite quadratic form by taking the absolute values of ρ_1, ρ_2. The matrix Q is obtained by introducing the linear parametrization of the surface $x = \sum_{i=1}^n \Phi_i(u) v_i$ into distance approximation (4) and the tangent terms.

Incorporating this approximation into the upper envelop corresponds to using a quasi-Newton algorithm optimizing jointly the $v_1 ... v_n$ and $u_1, ..., u_N$ [19, 30] but let us work with approximate correspondences. We update the surface parameters $u_1, \dots, u_N$ where we approximate each $\mathcal{D}_S^2(p_j)$ by sampling the current surface uniformly and creating a Kd-tree to find the sample closest to each p_j. This strategy also determines the parameters $\bar{u}_1, \dots, \bar{u}_m \in \partial\mathcal{T}_\mathbf{M}$ associated with boundary points $\bar{p}_1, \dots, \bar{p}_j$ to make the constraints $\bar{p}_j = \sum_{i=1}^n \Phi_i(\bar{u}_j) v_i$ linear in the optimization variables v and define matrix C.

Our algorithm avoids the slow convergence of coordinate descent and the large optimization problem of LM solvers [4, 11, 12]. An iteration of our algorithm solves a $3n \times 3n$ linear system while an LM iteration solves a $(3n{+}3N) \times (3n{+}3N)$ one with $N \gg n$. Comparing the convergence rate of LM and MM solvers for non-convex problems offers limited insight because the algorithms converge to different solutions and the rates are computed w.r.t. different optima.

7. Smooth Shape Analysis

We explain how to do smooth shape analysis with a subdivision surface. We focus on techniques invariant to isometries based on the Laplace-Beltrami operator.

Laplace-Beltrami Given a smooth shape S, the Laplace-Beltrami-Operator Δ maps a Sobolev functions $g \in H^1(S)$ to a continuous function $\Delta(g)\colon S \to \mathbb{R}$. Stokes' theorem $\int_S \Delta g \cdot h = \int_S \langle \nabla g, \nabla h \rangle$ let us discretize Δ in the space $H_n^1(S)$ spanned by Sobolev functions $\Phi_1 ... \Phi_n \in H^1(S)$ in terms of mass and stiffness matrices $D_0, D_1 \in \mathbb{R}^{n \times n}$

$$(D_0)_{ij} = \int_S \Phi_i \cdot \Phi_j \,\mathrm{dx} \quad (D_1)_{ij} = \int_S \langle \nabla \Phi_i, \nabla \Phi_j \rangle \,\mathrm{dx}.$$

that compute the scalar product of base functions and their tangent fields in S. The discrete Laplace-Beltrami operator in $H_n^1(S)$ is the linear operator $D_2 := D_0^{-1} D_1$ that has a non-negative spectrum and whose eigenvectors discretizes the eigenfunctions of Δ in $H_n^1(S)$. If $v \in \mathbb{R}^n$ is a D_2-eigenvector with eigenvalue λ, then $g_v = \sum_{i=1}^n v_i \Phi_i$ is a Δ-eigenfunction with eigenvalue λ in $H_n^1(S)$.

To particularize this construction to subdivision surfaces S, we must parametrize Sobolev functions over the surface $g \in H^1(S)$ by n scalar values $g_1 ... g_n \in \mathbb{R}$ with a mapping

$$g\colon S \to \mathbb{R} \tag{5}$$

$$\sum_{i=1}^n \Phi_i(u, f) \cdot v_i \mapsto \sum_{i=1}^n \Phi_i(u, f) \cdot g_i \tag{6}$$

that depends on each basis function Φ_i twice, once to describe the function g and once to describe the function domain S. As a result, the gradient of a surface function is

$$\nabla g(X(u)) = JX(u) G^{-1}(u) \sum_{i=1}^n \nabla \Phi_i(u) g_i,$$

where $JX(u)$ is the Jacobian of the surface mapping X with respect to u, and $G(u) = JX(u)^\top JX(u)$ is the first fundamental form of S at surface point $X(u)$. We compute the surface integrals in the mass and stiffness matrices by pulling-back the integrand to the quad mesh and using 3×3 Gaussian quadrature over each quad $\square$, that is

$$\int_S g(s) \,\mathrm{ds} = \sum_{f \in F} \int_\square g \circ X(u) \sqrt{\det G(u)} \,\mathrm{du}.$$

Geodesic Distances To compute the geodesic distance χ to a surface point $S(u_0)$ we need to solve the Eikonal equation $\|\nabla \chi\| = 1$ subject to boundary conditions $\chi(u_0) = 0$. As this nonlinear PDE is expensive to solve, [6] computes an approximate geodesic by first computing a function $f \in H^1(S)$ whose gradient is parallel to $\nabla\chi$, approximating $\nabla\chi$ by the unit vector field $X = -\nabla f / / \|\nabla f\|$, and finding the minimizer of $\int_S \|\nabla\chi - X\|^2$. As all the computations only require the Heat kernel Δ, we can use this method to directly approximate geodesics with the Laplace-Beltrami-Operator.

8. Experiments

We use point clouds acquired with a range sensor in a realistic setting for shape analysis, with partial and full object scans obtained by estimating the camera motion with off-the-shelf software and projecting the depth into a reference frame. Points within a bounding box are added to the point cloud that has the levels of noise, holes, and misalignment characteristic of computer vision systems. We also include experiments fitting 72 artifact-free high-resolution meshes from TOSCA for quantitative comparisons. TOSCA has high and low-resolution versions of meshes that let us to compare our subdivision representation to high and low-resolution triangular meshes of a shape, their descriptors, geodesics, and matchings. In our experiments, we coarsely choose the model parameters for each shape and method and answer the following questions: Our technical report [7] includes additional experiments.

How does our subdivision surface fit compare to triangular meshes? Our technique fits a surface to a noisy pointcloud as accurately as the mesh obtained by Poisson reconstruction [14] with 2% of its vertices (Fig. 1, 3, 9).

How does our fitting technique compare to subdivision state-of-the art? We fit a subdivision surface with 1000-2000 control vertices to high-resolution meshes and measure the distance (Table 1) between the mesh and the fitted subdivision surfaces with our method and the techniques [4, 12, 22, 29] closer it. **In terms of accuracy**, our model is more robust to noise than the least-squares model of [22] and avoids the artifacts caused by noise and outliers (Fig. 3 and 10). It also fits better small curved structures in clean meshes (Fig. 5) because our second-order distance approximation incorporates curvature. **In terms of optimization**, our algorithm is an order of magnitude faster and

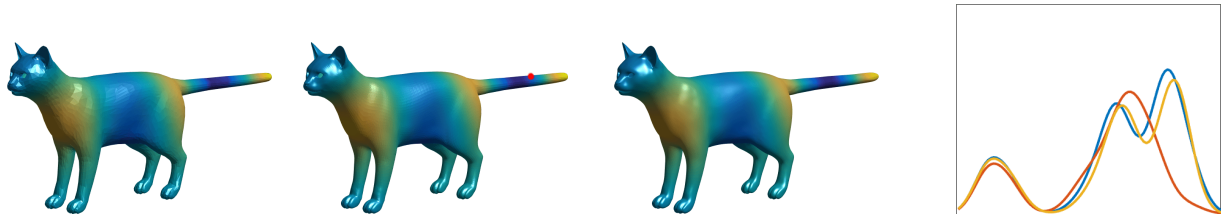

Mesh 3,400 vertices | Mesh 27,984 vertices | Subdiv. Surf. 1700 vertices | wave-kernel signatures (WKS)

Figure 2. 12^{th} Δ-eigenfunction and WKS computed at the point in red with fine and coarse triangular meshes (blue and red lines) and a subdivision surface (yellow). The WKS of the subdivision surface and fine mesh are similar, the coarse mesh merges high-frequency peaks.

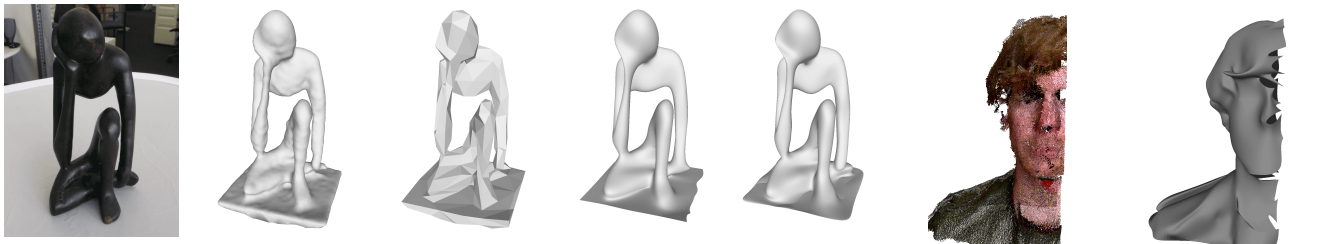

object 110 K points | mesh by [14] $|V| = 20K$ | collapsed mesh $|V| = 250$ | [22] SDS $|V| = 250$ | our SDS $|V| = 250$ | pointcloud | [22] subdiv. surf.

pointcloud 140 K points | mesh by [14] $|V| = 12K$ | collapsed mesh $|V| = 250$ | [22] SDS $|V| = 250$ | our SDS $|V| = 250$ | LM solver | our constrained model

Figure 3. Surfaces reconstructed from Kinect points with Poisson method [14] and subdivision surface (SDS) methods. Our technique reconstructs the surface with an accuracy comparable to [14] high-resolution mesh with a parametrization that has 2% of the vertices (V). An edge-collapsed mesh of comparable size looses all small-scale details. The subdivision surface of [22] shrinks the statue's V-shaped torso, oversmooths its hands, and produces artifacts at the base of the bottle where samples are irregularly spaced and overlap.

Figure 4. Subdiv. surfaces with 500 control vertices fitted to a partial Kinect scan (260K points). Our model constrains the boundaries of the surface to match the boundary points, while the unconstrained models [22] and the LM [4, 11, 12, 29] shrink the boundaries by regularization.

less sensitive to local minima than using the LM solver advocated by [4, 11, 12, 29] to minimize the MM envelope. In particular, our surfaces evolve further from initialization and reproduce the small-scale details of the input meshes (Figs. 5 and 10). Our technique is moderately slower than the quadratic model of [22] due to the robust energy. Table 1 show the effects of each term in our energy model.

Are shape descriptors obtained with subdivision discriminative? The wave kernel signature (WKS) [1] assigns to each point $x \in S$ a function w_x that depends on the value of the Δ-eigenfunctions at x. Figure 2 compares WKS obtained with different surface discretizations: a fine triangular mesh, a coarse triangular mesh, and our subdivision surface. The WKS of the fine mesh presents two frequency peaks that are preserved in the WKS of subdivision, while the WKS from the coarse mesh looses the high-frequency details.

Table 1. Average run time and Hausdorff distance (as ‰of the bounding box) between a 1K-vertex subdivision surface and high-resolution input meshes. Columns 1–3, 6–8 show how each additional term in the energy slows the optimization but improves the accuracy ($\mathcal{D}^2$ penalizes the squared distance instead of $\mathcal{D}$).

	execution time (s)					Hausdorff Distance				
mesh	$\mathcal{D}^2$	$\alpha = 0$	our	[22]	LM	$\mathcal{D}^2$	$\alpha = 0$	our	[22]	LM
cat	27	37	44	17	**384**	0.28	0.26	**0.24**	0.32	0.37
centaur	16	18	24	8	**228**	0.52	0.47	**0.45**	0.53	0.61
david	58	51	82	36	**607**	0.37	0.34	**0.33**	0.38	0.56
dog	24	28	40	15	**304**	0.39	0.36	**0.35**	0.44	0.51
gorilla	38	41	57	25	**513**	0.71	0.57	**0.57**	0.70	0.99
horse	21	25	34	12	**246**	0.36	0.34	**0.34**	0.40	0.65
michael	58	59	89	35	**535**	0.43	0.38	**0.36**	0.43	0.57
victoria	44	505	59	26	**423**	0.33	0.29	**0.29**	0.35	0.51
wolf	6	9	10	6	**65**	0.42	0.43	**0.41**	0.47	0.53

Can we do shape matching with subdivision representations? Functional maps [24] formulate the problem

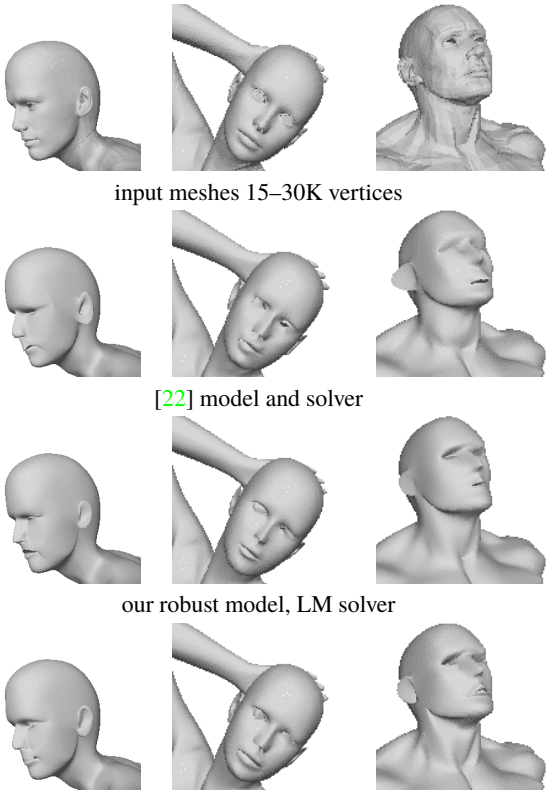

Figure 5. Comparison of subdivision surface techniques. All the subdivision surfaces have 1700 control vertices and are visualized by refining $\mathbf{M}_0$ 3 times zooming into the detailed surface face.

Figure 6. Fitted colored subdiv. surfaces by 2K control parameters.

of matching points in two shapes as matching smooth functions parametrized by Δ-eigenfunctions over them. Fig. 8 shows the computed matching with respect to two different shape representations: a fine mesh with 27894 vertices and a subdivision surface with 1700 control vertices. The matching results are comparable (small accuracy loss at the tails of some cats) even though the surface representation is 10 times smaller. In Fig. 8 we color the template shape with a map determined by the spatial coordinates of every point and transfer this color map to the target shape with the estimated matching. Matching with a coarse mesh fails.

Are geodesics computed with subdivision accurate? Fig. 1 and 7 compare geodesics computed by the Heat

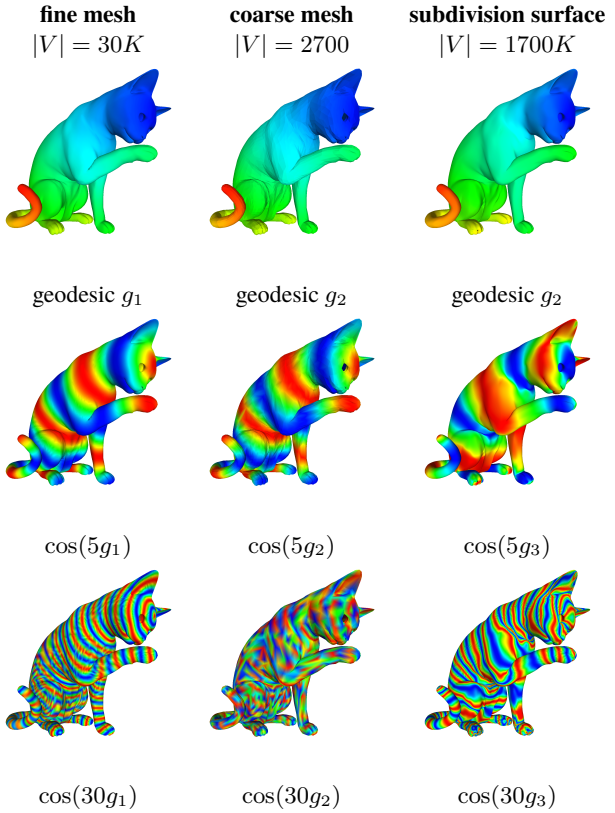

Figure 7. Approximate geodesics [6] computed with different surface representations: a high-resolution mesh, a low-resolution mesh, and our subdivision surface. The level lines of the geodesics g are visualized as $\cos(\varpi g)$ to show that our subdivision surface is comparable to the fine mesh and can represent geodesics at both low and high resolution with a parametrization with 5% of vertices (V). In contrast, the low-resolution mesh looses all the small scale details of the surface and high-frequency geodesic information.

method [6] with different surface representations. As ground-truth proxy, we use a high-resolution mesh obtained by the Poisson reconstruction [14] and compare it to geodesics computed over low-resolution meshes and our subdivision representation. The geodesics computed with subdivision are comparable to ones computed with high-resolution meshes, while the geodesics of low-resolution meshes cannot represent the high-frequency information.

9. Conclusions

We have presented a method to fit a subdivision surface to low-level surface representations and develop standard shape analysis with it. Our fitting model is robust to noise and outliers and the resulting subdivision surfaces represent smooth shapes at high accuracy with a fraction of the variables used with triangular meshes. Our experiments show how the smoothness of the representation reduces the size of the Laplace-Beltrami operator without significant loss of accuracy and let us compute shape descriptors, approximate geodesics, and shape matches compactly.

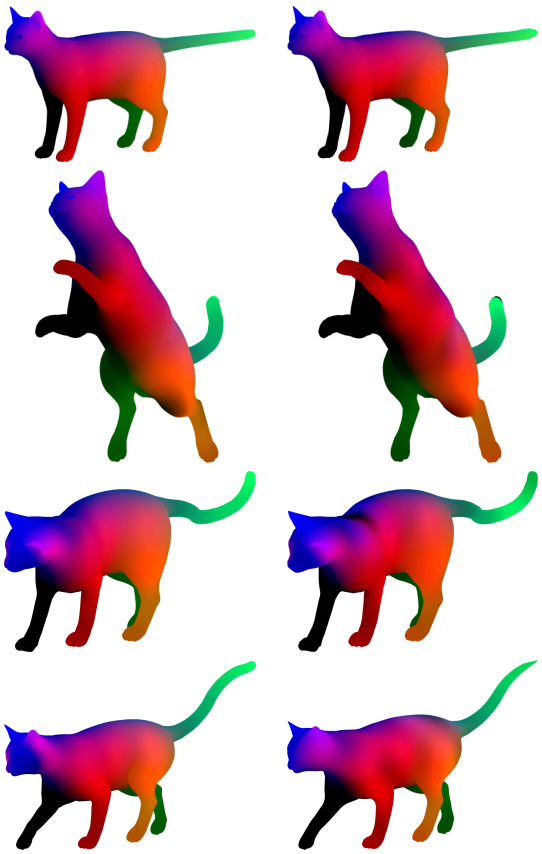

Figure 8. Shape matching by functional maps [24] with subdivision representations. The color of the target shape (left) is transferred to the matched shape (right) with the estimated functional mapping.

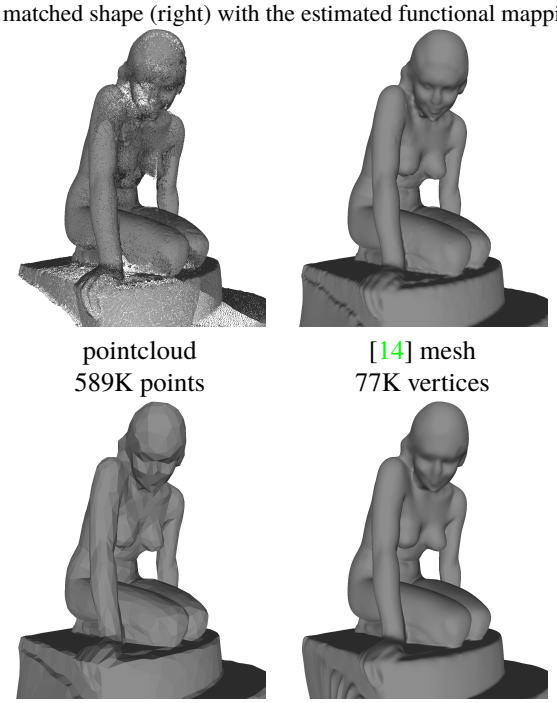

pointcloud
589K points

[14] mesh
77K vertices

collapsed mesh
2K vertices

our subdiv. surface
2K vertices

Figure 9. Surfaces reconstructed from a pointcloud with the Poisson surface reconstruction method [14] and with our subdivision surface technique. The accuracy of our subdivision surface is comparable to the high-resolution mesh but has 5% of the mesh vertices. An edge-collapsed mesh of comparable size is unable to represent smooth textures and fine details (belly bottom).

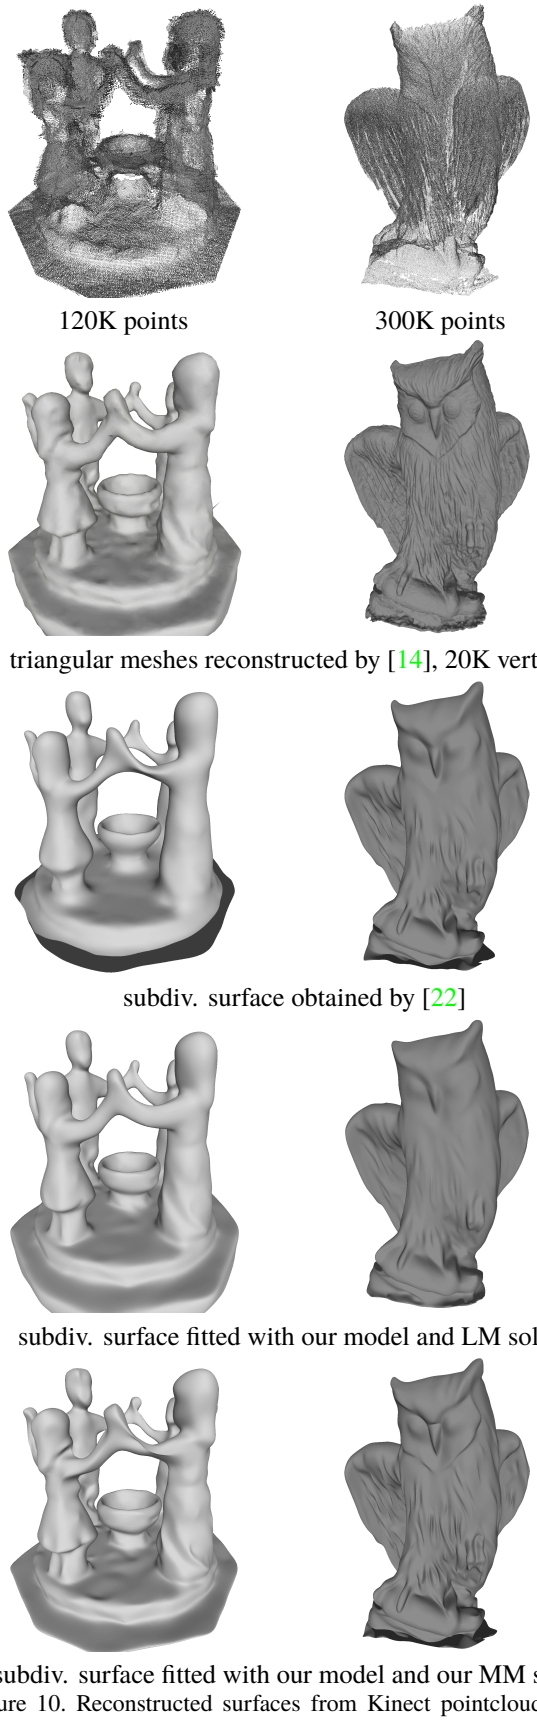

120K points

300K points

triangular meshes reconstructed by [14], 20K vertices

subdiv. surface obtained by [22]

subdiv. surface fitted with our model and LM solver

subdiv. surface fitted with our model and our MM solver

Figure 10. Reconstructed surfaces from Kinect pointclouds by the Poisson method [14] and different subdivision surfaces of 2000 control vertices. Our model reconstructs the surface with an accuracy comparable to the Poisson mesh with a parametrization 20 times smaller. The least-squares model [22] shrinks the dancer's arms and oversmooths the owls neck, while the LM solver [4, 11, 12, 29] optimizes our model to a lower accuracy than our MM algorithm (owls wings).

References

[1] M. Aubry, U. Schlickewei, and D. Cremers. The wave kernel signature: A quantum mechanical approach to shape analysis. In *IEEE ICCV*, 1626–1633, 2011. 4321, 4326

[2] A. Bronstein, M. Bronstein, and R. Kimmel. *Numerical Geometry of Non-Rigid Shapes*. Springer, 2008. 4321

[3] D. Burkhart, B. Hamann, and G. Umlauf. Isogeometric finite element analysis based on Catmull-Clark subdivision solids. *Eurographics Symp. Geom. Process.*, 29(5):1575–1584, 2010. 4322

[4] T. J. Cashman and A. W. Fitzgibbon. What shape are dolphins? building 3D morphable models from 2D images. *IEEE Trans. Pattern Anal. Mach. Intell.*, 35(1):232–244, 2013. 4321, 4322, 4324, 4325, 4326, 4328

[5] E. Catmull and J. Clark. Recursively generated b-spline surfaces on arbitrary topological meshes. *Computer-Aided Design*, 10(6):350 – 355, 1978. 4323

[6] K. Crane, C. Weischedel, and M. Wardetzky. Geodesics in heat: A new approach to computing distance based on heat flow. *ACM Trans. Graph.*, 32(3):10, 2013. 4321, 4322, 4325, 4327

[7] V. Estellers, F. R. Schmidt, and D. Cremers. Compression for Smooth Shape Analysis. *arXiv.*, 1711.10824, 2017. 4325

[8] M. Garland and P. S. Heckbert. Surface simplification using quadric error metrics. *ACM Trans. Graph SIGGRAPH*, 209–216, 1997. 4322, 4324

[9] F. Goes, M. Desbrun, M. Meyer, and T. Derose. Subdivision Exterior Calculus for Geometry Processing. *ACM Trans. Graph.*, 35(10), 2016. 4322

[10] H. Hoppe, T. DeRose, T. Duchamp, M. Halstead, H. Jin, J. McDonald, J. Schweitzer, and W. Struetzle. Piecewise smooth surface reconstruction. *ACM Comput. Graph.*, 294–302, 1994. 4321

[11] S. Ilic. Using Subdivision Surfaces for 3D Reconstruction from Noisy Data. *Deform. Image Regist. Deform. Environ.*, 2006. 4321, 4322, 4324, 4325, 4326, 4328

[12] M. Jaimez, T. J. Cashman, A. Fitzgibbon, J. Gonzalez-jimenez, and D. Cremers. An Efficient Background Term for 3D Reconstruction and Tracking with Smooth Surface Models. *IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017. 4321, 4322, 4324, 4325, 4326, 4328

[13] B. Jüttler, A. Mantzaflaris, R. Perl, and M. Rumpf. On numerical integration in isogeometric subdivision methods for PDEs on surfaces. *Computer Methods in Applied Mechanics and Engineering*, 302:131 – 146, 2016. 4322

[14] M. Kazhdan and H. Hoppe. Screened poisson surface reconstruction. *ACM Trans. Graph.*, 32(3), 2013. 4322, 4324, 4325, 4326, 4327, 4328

[15] V. Kolmogorov. Blossom V: a new implementation of a minimum cost perfect matching algorithm. *Mathematical Programming Computation*, 1(1):43–67, 2009. 4324

[16] G. Lavoué, F. Dupont, and A. Baskurt. Subdivision surface fitting for efficient compression and coding of 3D models. *Vis. Commun. Image Process.*, 5960(33):59603F–59603F–12, 2005. 4321, 4322

[17] B. Levy. Laplace-Beltrami Eigenfunctions: Towards an Algorithm that Understands Geometry. *IEEE Int. Conf. Shape Model. Appl.*, 2006. 4321

[18] R. Litman and A. M. Bronstein. Learning spectral descriptors for deformable shape correspondence. *IEEE Trans. Pattern Anal. Mach. Intell.*, 36(1):171–180, 2014. 4321

[19] Y. Liu and W. Wang. A revisit to least squares orthogonal distance fitting of parametric curves and surfaces. *Lect. Notes Comput. Sci.*, 4975 LNCS(1):384–397, 2008. 4322, 4325

[20] W. E. Lorensen and H. E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. *ACM SIGGRAPH Comput. Graph.*, 21(4):163–169, 1987. 4324

[21] W. Ma, X. Ma, S. K. Tso, and Z. Pan. Subdivision surface fitting from a dense triangle mesh. *Proc. - Geom. Model. Process. Theory Appl. GMP 2002*, 36:94–103, 2002. 4321

[22] M. Marinov and L. Kobbelt. Optimization methods for scattered data approximation with subdivision surfaces. *Graph. Models*, 67(5):452–473, 2005. 4321, 4322, 4325, 4326, 4327, 4328

[23] D. Zorin. Subdivision for Modeling and Animation. *New York*, 98:1 – 203, 2000. 4321, 4323

[24] M. Ovsjanikov, M. Ben-Chen, J. Solomon, A. Butscher, and L. Guibas. Functional maps. *ACM Trans. Graph.*, 31(4), 2012. 4321, 4326, 4328

[25] H. Pottmann and S. Leopoldseder. A concept for parametric surface fitting which avoids the parametrization problem. *Comput. Aided Geom. Des.*, 20(6):343–362, 2003. 4324

[26] J. F. Remacle, J. Lambrechts, B. Seny, E. Marchandise, A. Johnen, and C. Geuzainet. Blossom-Quad: A non-uniform quadrilateral mesh generator using a minimum-cost perfect-matching algorithm. *Int. J. Numer. Methods Eng.*, 89(9):1102–1119, 2012. 4324

[27] M. Reuter, F.-E. Wolter, and N. Peinecke. Laplace-Beltrami Spectra As Shape-DNA of Surfaces and Solids. *Comput. Aided Des.*, 38(4):342–366, 2006. 4321

[28] J. Stam. Exact evaluation of Catmull-Clark subdivision surfaces at arbitrary parameter values. *ACM Trans. Graph. SIGGRAPH*, 395–404, 1998. 4323

[29] J. Taylor, J. Valentin, B. Luff, R. Banks, A. Fitzgibbon, T. Sharp, E. Wood, and J. Shotton. Efficient and Precise Interactive Hand Tracking Through Joint , Continuous Optimization of Pose and Correspondences. *ACM Trans. Graph. SIGGRAPH*, 35(4):143, 2016. 4321, 4322, 4325, 4326, 4328

[30] W. Wang, H. Pottmann, and Y. Liu. Fitting B-spline curves to point clouds by curvature-based squared distance minimization. *ACM Trans. Graph.*, 25(2):214–238, 2006. 4322, 4325

[31] A. Wawrzinek and K. Polthier. Computer-Aided Design Integration of generalized B-spline functions on Catmull-Clark. *Comput. Des.*, 78:60–70, 2016. 4322

[32] T. Windheuser, M. Vestner, E. Rodola, R. Triebel, and D. Cremers. Optimal Intrinsic Descriptors for Non-Rigid Shape Analysis. *Br. Mach. Vis. Conf.*, 2014. 4321

[33] W. Zheng, P. Bo, Y. Liu, and W. Wang. Fast B-spline curve fitting by L-BFGS. *Comput. Aided Geom. Des.*, 29(7):448–462, 2012. 4322