

Static/Dynamic Filtering for Mesh Geometry

Juyong Zhang, Bailin Deng[†], Yang Hong, Yue Peng, Wenjie Qin, and Ligang Liu

Abstract—The joint bilateral filter, which enables feature-preserving signal smoothing according to the structural information from a guidance, has been applied for various tasks in geometry processing. Existing methods either rely on a static guidance that may be inconsistent with the input and lead to unsatisfactory results, or a dynamic guidance that is automatically updated but sensitive to noises and outliers. Inspired by recent advances in image filtering, we propose a new geometry filtering technique called *static/dynamic filter*, which utilizes both static and dynamic guidances to achieve state-of-the-art results. The proposed filter is based on a nonlinear optimization that enforces smoothness of the signal while preserving variations that correspond to features of certain scales. We develop an efficient iterative solver for the problem, which unifies existing filters that are based on static or dynamic guidances. The filter can be applied to mesh face normals followed by vertex position update, to achieve scale-aware and feature-preserving filtering of mesh geometry. It also works well for other types of signals defined on mesh surfaces, such as texture colors. Extensive experimental results demonstrate the effectiveness of the proposed filter for various geometry processing applications such as mesh denoising, geometry feature enhancement, and texture color filtering.

Index Terms—Geometry Processing, Mesh Filtering, Mesh Denoising.



1 INTRODUCTION

SIGNAL filtering, the process of modifying signals to achieve desirable properties, has become a fundamental tool for different application areas. In image processing, for example, various filters have been developed for smoothing images while preserving sharp edges. Among them, the *bilateral filter* [1] updates an image pixel using the weighted average of nearby pixels, taking into account their spatial and range differences. Its simplicity and effectiveness makes it popular in image processing, and inspires various follow-up work with improved performance [2], [3], [4], [5].

Besides image processing, filtering techniques have also been utilized for processing 3D geometry. Indeed, many geometric descriptors such as normals and vertex positions can be considered as signals defined on two-dimensional manifold surfaces, where image filtering methods can be naturally extended and applied. For example, the bilateral filter has been adapted for feature-preserving mesh smoothing and denoising [6], [7], [8], [9].

Development of new geometry filters has also been inspired by other techniques that improve upon the original bilateral filter. Among them, the *joint bilateral filter* [2], [3] determines the filtering weights using the information from a *guidance* image instead of the input image, and achieves more robust filtering results when the guidance provides reliable structural information. One limitation of this approach is that the guidance image has to be specified beforehand, and remains *static* during the filtering processing. For image texture filtering, Cho et al. [4] address this issue by computing the guidance using a patch-based approach that reliably captures the image structure. This idea was

later adopted by Zhang et al. [10] for mesh denoising, where a patch-based guidance is computed for filtering the face normals. Another improvement for the joint bilateral filter is the *rolling guidance filter* proposed by [5], which iteratively updates an image using the previous iterate as a *dynamic* guidance, and is able to separate signals at different scales. Recently, this approach was adapted by Wang et al. [11] to derive a *rolling guidance normal filter* (RGNF), with impressive results for scale-aware geometric processing.

For guided filtering, the use of static vs dynamic guidance presents a trade-off between their properties. Static guidance enables direct and intuitive control over the filtering process, but is not trivial to construct a priori for general shapes. Dynamic guidance, such as the one used in RGNF, is automatically updated according to the current signal values, but can be less robust when there are outliers or noises in the input signal. Recently, Ham et al. [12] combine static and dynamic guidance for robust image filtering. Inspired by their work, we propose in this paper a new approach for filtering signals defined on mesh surfaces, by utilizing both static and dynamic guidances. The filtered signal is computed by minimizing a target function that enforces consistency of signal values within each neighborhood, while incorporating structural information provided by a static guidance. To solve the resulting nonconvex optimization problem, we develop an efficient fixed-point iteration solver, which significantly outperforms the majorization-minimization (MM) algorithm proposed by [12] for similar problems. Moreover, unlike the MM algorithm, our solver can handle constraints such as unit length for face normals, which are important for geometry processing problems. Our solver iteratively updates the signal values by combining the original signal with the current signal from a spatial neighborhood. The combination weights are determined according to the static input guidance as well as a dynamic guidance derived from the current signal. The proposed method, called *static/dynamic* (SD) filtering, benefits from both types of guidance and produces

• J. Zhang, Y. Hong, Y. Peng, W. Qin, and L. Liu are with School of Mathematical Sciences, University of Science and Technology of China.

• B. Deng is with School of Computer Science and Informatics, Cardiff University.

[†]Corresponding author. Email: DengB3@cardiff.ac.uk.

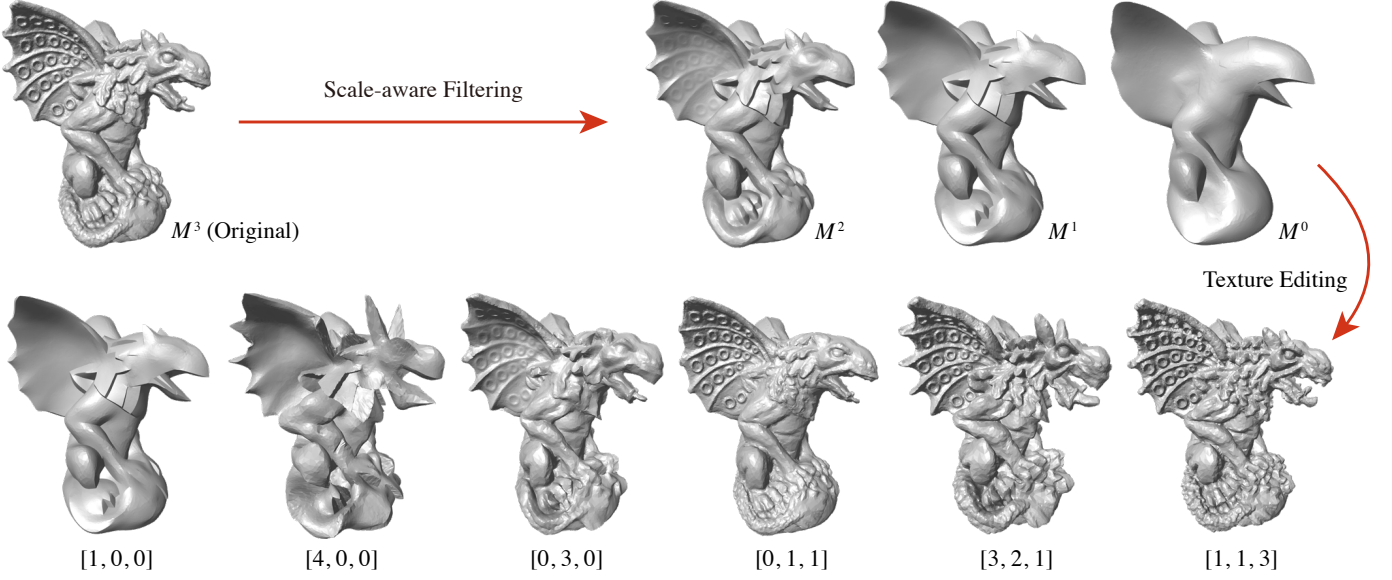


Fig. 1. Our SD filter can be used for scale-aware filtering of mesh geometry, allowing us to separate geometry signals according to their scales. Such decomposition can be used for manipulating geometric details, boosting or attenuating features at different scales.

scale-aware and feature-preserving results.

The proposed method can be applied to different signals on mesh surfaces. When applied to face normals followed by vertex updates, it filters geometric features according to their scales. When applied to mesh colors obtained from texture mapping, it filters the colors based on the metric on the mesh surface. In addition, utilizing the scale-awareness of the filter, we apply it repeatedly to separate signal components of different scales; the results can be combined according to user-specified weights, allowing for intuitive feature manipulation and enhancement. Extensive experimental results demonstrate the efficiency and effectiveness of our filter. We also release the source codes to ensure reproducibility.

In addition, we propose a new method for vertex update according to face normals, using a nonlinear optimization formulation that enforces the face normal conditions while preserving local triangle shapes. The vertex positions are computed by iteratively solving a linear system with a fixed sparse positive definite matrix, which is done efficiently via pre-factorization of the matrix. Compared with existing approaches, our method produces meshes that are more consistent with the filtered face normals.

In summary, our main contributions include:

- we extend the work of Ham et al. [12] and propose an SD filter for signals defined on triangular meshes, formulated as an optimization problem;
- we develop an efficient fixed-point iteration solver for the SD filter, which can handle constraints such as unit normals and significantly outperforms the MM solver from [12];
- we propose an efficient approach for updating vertex positions according to filtered face normals, which produces new meshes that are consistent with the target normals while preserving local triangle shapes;
- based on the SD filter, we develop a method to separate and combine signal components of different scales, enabling intuitive feature manipulation for mesh geometry and texture color.

2 RELATED WORK

In the past, various filtering approaches have been proposed to process mesh geometry. Early work from Taubin [13] and Desbrun et al. [14] applied low-pass filters on meshes, which remove high-frequency noises but also attenuate sharp features. Later, Taubin [15] proposed a two-step approach that first performs smoothing on face normals, followed by vertex position updates using anisotropic filters. To enhance crease edges, Ohtake et al. [16] applied anisotropic diffusion to mesh normals before updating vertex positions. Chuang and Kazhdan [17] developed a framework for curvature-aware mesh filtering based on the screened Poisson equation.

An important class of mesh filtering techniques is based on the bilateral filter [1]. On images, the bilateral filter updates a pixel using a weighted average of its neighboring pixels, with larger contribution from pixels that are closer in spatial or range domain. It can smooth images while preserving edges where there is large difference between neighboring pixel values [18]. Different methods have been developed to adapt the bilateral filter to mesh geometry. Fleishman et al. [6] and Jones et al. [7] applied the bilateral filter to the mesh vertex positions for feature-preserving mesh denoising. Zheng et al. [8] applied the bilateral filter to mesh face normals instead, followed by vertex position update to reconstruct the mesh shape. Solomon et al. [9] proposed a framework for bilateral filter that is applicable for signals on general domains including images and meshes, with a rigorous theoretical foundation. Besides denoising, bilateral filtering has also been applied for other geometry processing applications such as point cloud normal enhancement [19] and mesh feature recovery [20].

The bilateral filter inspired a large amount of follow-up work on image filtering. Among them, the joint bilateral filter [2], [3] extends the original bilateral filter by evaluating the spatial kernel using a guidance image. It can produce more reliable results when the guidance image correctly captures the structural information of the target signal.

This property was utilized by Eisemann & Durand [2] and Petschnigg et al. [3] to filter flash photos, using corresponding non-flash photos as the guidance. Kopf et al. [21] and Cho et al. [4] applied the joint bilateral filter for image upsampling and structure-preserving image decomposition, respectively. In particular, a patch-based guidance is constructed in [4] to capture the input image structure. This idea was later adopted in [10] for filtering mesh face normals, where the guidance normals are computed using surface patches with the most consistent normals. Zhang et al. [5] proposed a different approach to guidance construction in their iterative *rolling guidance filter*, where the resulting image from an iteration is used as a dynamic guidance for the next iteration. The rolling guidance filter produces impressive results for scale-aware image processing, and is able to filter out features according to their scales. Wang et al. [11] adapted this approach to filter mesh face normals; the resulting *rolling guidance normal filter* enables scale-aware processing of geometric features, but is sensitive to noises on the input model. Recently, Ham et al. [12] proposed a robust image filtering technique based on an optimization formulation that involves a nonconvex regularizer. Their technique is effectively an iterative filter that incorporates both static and dynamic guidances, and achieves superior results in terms of robustness, feature-preservation, and scale-awareness. Our SD filter is based on a similar optimization formulation, but takes into account the larger filtering neighborhoods that are necessary for geometry signals. It enjoys the same desirable properties as its counterpart in image processing. In addition, the numerical solver proposed in [12] can only handle unconstrained signals, and is less efficient for the large neighborhoods used in our formulation. We therefore propose a new solver that outperforms the one from [12], while allowing for constrained signals such as unit normals.

Feature-preserving signal smoothing can also be achieved via optimization. Notable examples include image smoothing algorithms that induce sparsity of image gradients via ℓ_0 -norm [22] or ℓ_1 -norm [23] regularization. These approaches were later adapted for mesh smoothing and denoising [24], [25], [26]. Although effective in many cases, their optimization formulation only regularizes the signal difference between immediately neighboring faces. In comparison, our optimization compares signals within a neighborhood with user-specified size, which provides more flexibility and achieves better preservation of large-scale features.

From a signal processing point of view, meshes can be seen as a combination of signals with multiple frequency bands, which also relates with the scale space analysis [27]. Previous work separate geometry signals of different frequencies using eigenfunctions of the heat kernel [28] or the Laplace operator [29], [30]. Although developed with sound theoretical foundations, such approaches are computationally expensive. Moreover, as specific geometric features can span across a wide range of frequencies, it is not easy to preserve or manipulate them with such approaches. The recent work from Wang et al. [11] provides an efficient way to separate and edit geometric features of different scales, harnessing the scale-aware property of the rolling guidance filter. Our SD filter also supports scale-aware processing of geometry signals, with more robustness than RGNF thanks to the incorporation of both static and dynamic guidances.

3 THE SD FILTER

The SD filter was originally proposed by Ham et al. [12] for robust image processing. Given an input image F and a static guidance image G , they compute an output image U via optimization

$$\min_U \sum_i \gamma_i (U_i - F_i)^2 + \lambda \sum_{(i,j) \in \mathcal{N}} \phi_\mu(G_i - G_j) \cdot \psi_\nu(U_i - U_j), \quad (1)$$

where F_i, G_i, U_i are the pixel values of F, G and U respectively, γ_i and λ are user-specified weights, \mathcal{N} denotes the set of 8-connected neighboring pixels, and

$$\phi_\mu(x) = \exp\left(-\frac{x^2}{2\mu^2}\right), \quad \psi_\nu(x) = 1 - \phi_\nu(x). \quad (2)$$

The first term in the target function is a fidelity term that requires the output image to be close to the input image, while the second term is a regularizer for the output image. Function ψ_ν (see Fig. 2) penalizes the difference between adjacent pixels, but with bounded penalty for pixel pairs with large difference which correspond to edges or outliers. When ν approaches 0, ψ_ν approaches the ℓ_0 norm. Function ϕ_μ is a Gaussian weight according to the guidance, with larger weights for pixel pairs with closer guidances. Thus the regularizer promotes smooth regions and preserves sharp features based on the guidance, and is robust to outliers.




Fig. 2. The graphs of $\psi_\nu(x)$ with different ν parameters.

In this paper, we propose an SD filter for signals defined on 2-manifold surfaces represented as triangular meshes. We begin our discussion with filtering face normals, a common approach for smoothing mesh geometry [8], [9], [10], [31].

3.1 SD filter for face normals

For a given orientable triangular mesh, let $\mathbf{n}_i \in \mathbb{R}^3$ be the outward unit normal of face f_i , computed as

$$\mathbf{n}_i = \frac{(\mathbf{v}_{i_1} - \mathbf{v}_{i_2}) \times (\mathbf{v}_{i_3} - \mathbf{v}_{i_2})}{\|(\mathbf{v}_{i_1} - \mathbf{v}_{i_2}) \times (\mathbf{v}_{i_3} - \mathbf{v}_{i_2})\|}, \quad (3)$$

where $\mathbf{v}_{i_1}, \mathbf{v}_{i_2}, \mathbf{v}_{i_3} \in \mathbb{R}^3$ are its vertex positions enumerated according to the orientation. We associate the normal with the face centroid $\mathbf{c}_i = \frac{1}{3} \sum_{k=1}^3 \mathbf{v}_{i_k}$. We first filter the face normals, and then update the mesh vertices accordingly. To define an SD filter for the normals $\{\mathbf{n}_i\}$, we must consider some major differences compared with image filtering:

- Image pixels are located on a regular grid, but mesh faces may result from irregular sampling of the surface.
- To smooth an image, the SD filter as per Eq. (1) only considers the difference between a pixel and its eight neighbor

pixels. On meshes, however, geometry features can span across a large region, thus we may need to compare face normals beyond one-ring neighborhoods [11]. Moreover, similar to the bilateral filter, such comparison should consider the difference between the spatial locations, with stronger penalty for normal deviation between faces that are closer to each other.

Therefore, we compute the filtered normals $\{\mathbf{n}_i\}$ by minimizing a target function

$$E_{SD} = E_{fid} + \lambda E_{reg}, \quad (4)$$

with a user-specified weight $\lambda > 0$. Here E_{fid} is a fidelity term between the input and output normals,

$$E_{fid} = \sum_i A_i \|\mathbf{n}_i - \hat{\mathbf{n}}_i\|^2, \quad (5)$$

where $\hat{\mathbf{n}}_i, A_i$ are the normal and area of face f_i , respectively. E_{reg} is a regularization term defined as

$$E_{reg} = \sum_i \sum_{f_j \in N(i)} [A_j \cdot \phi_\eta(\|\mathbf{c}_i - \mathbf{c}_j\|) \cdot \phi_\mu(\|\mathbf{g}_i - \mathbf{g}_j\|) \cdot \psi_\nu(\|\mathbf{n}_i - \mathbf{n}_j\|)], \quad (6)$$

where $\{\mathbf{g}_i\}$ are the guidance face normals, and $N(i)$ denotes the set of neighboring faces of f_i . The Gaussian standard deviation parameters $\eta, \mu, \nu \in \mathbb{R}^+$ are controlled by the user. Compared with the image regularizer in Eq. (1), this formulation introduces a Gaussian weight ϕ_η for the spatial locations of face normals. Here ϕ_η is defined according to the Euclidean distance between face centroids for simplicity of computation, but other distance measures such as the geodesic distance can also be used. For each face f_i , its neighborhood $N(i)$ is chosen to be the set of faces with a significant value of the spatial weight $\phi_\eta(\|\mathbf{c}_i - \mathbf{c}_j\|)$. Using the empirical three-sigma rule [32], we include in $N(i)$ the faces $\{f_j\}$ with $\|\mathbf{c}_j - \mathbf{c}_i\| \leq 3\eta$, which can be found using a breadth-first search from f_i .

The target function E_{SD} is nonconvex because of ψ_ν , and needs to be minimized numerically. In the following, we first show how the majorization-minimization (MM) algorithm proposed in [12] can be extended to solve this problem. Afterwards, we propose a new fixed-point iteration solver that significantly outperforms the MM algorithm and is suitable for interactive applications.

MM algorithm. For the SD image filter, Ham et al. [12] proposed a majorization-minimization (MM) algorithm to iteratively minimize the target function (1). In each iteration, the target function is replaced by a convex surrogate function that bounds it from above, which is computed using the current variable values. This surrogate function is then minimized to update the variables. The MM solver is guaranteed to converge to a local minimum of the target function. Thus a straightforward way to minimize the new target function (4) is to employ the MM algorithm, using the convex surrogate function Ψ_ν^t for $\psi_\nu(x)$ at $x = t$ [12]:

$$\Psi_\nu^t(x) = \psi_\nu(t) + (x^2 - t^2)(1/2\nu^2 - \psi_\nu(t)/2\nu^2). \quad (7)$$

Specifically, with the variable values $\{\mathbf{n}_i^k\}$ at iteration k , we replace the term $\psi_\nu(\|\mathbf{n}_i - \mathbf{n}_j\|)$ in the target function by its convex surrogate $\Psi_\nu^{\|\mathbf{n}_i^k - \mathbf{n}_j^k\|}(\|\mathbf{n}_i - \mathbf{n}_j\|)$ according to Eq. (7).

The updated variable values $\{\mathbf{n}_i^{k+1}\}$ are computed from the resulting convex problem

$$\min_{\{\mathbf{n}_i^{k+1}\}} \sum_i A_i \|\mathbf{n}_i^{k+1} - \hat{\mathbf{n}}_i\|^2 + \lambda \sum_i \sum_{f_j \in N(i)} w_{ij}^k \|\mathbf{n}_i^{k+1} - \mathbf{n}_j^{k+1}\|_2^2, \quad (8)$$

where

$$w_{ij}^k = \frac{A_j}{2\nu^2} \cdot \phi_\eta(\|\mathbf{c}_i - \mathbf{c}_j\|) \cdot \phi_\mu(\|\mathbf{g}_i - \mathbf{g}_j\|) \cdot \phi_\nu(\|\mathbf{n}_i^k - \mathbf{n}_j^k\|). \quad (9)$$

Due to the symmetry of neighboring relation between faces (i.e., $f_j \in N(i) \Leftrightarrow f_i \in N(j)$), the optimization problem (8) amounts to solving a linear system:

$$(\mathbf{D} + \lambda \mathbf{M}^k) \mathbf{N}^{k+1} = \mathbf{D} \hat{\mathbf{N}}, \quad (10)$$

where $\mathbf{D} = \text{diag}(A_1, A_2, \dots, A_{n_f})$ with n_f being the number of faces, $\mathbf{N}^{k+1}, \hat{\mathbf{N}} \in \mathbb{R}^{n_f \times 3}$ stack the values of $\{\mathbf{n}_i^{k+1}\}$ and $\{\hat{\mathbf{n}}_i\}$ respectively, and $\mathbf{M}^k \in \mathbb{R}^{n_f \times n_f}$ is a symmetric matrix with diagonal elements

$$m_{ii}^k = \sum_{j \in N(i)} (w_{ij}^k + w_{ji}^k),$$

and off-diagonal elements

$$m_{ij}^k = \begin{cases} -w_{ij}^k - w_{ji}^k, & \text{if } j \in N(i), \\ 0, & \text{otherwise.} \end{cases}$$

The linear system matrix in Eq. (10) is diagonally dominant and symmetric positive definite, and can be solved using standard linear algebra routines.

Fixed-point iteration solver. Although the MM algorithm works well on images, its performance on meshes is often unsatisfactory. Due to larger face neighborhoods, there are a large number of nonzeros in the linear system matrix of Eq. (10), resulting in long computation time for each iteration. In the following, we propose a more efficient solver that is suitable for interactive applications. Note that a local minimum of the target function (4) should satisfy the first order optimality condition $\partial E_{SD} / \partial \mathbf{n}_i = \mathbf{0}$ for each \mathbf{n}_i , which expands into

$$A_i(\mathbf{n}_i - \hat{\mathbf{n}}_i) + \lambda \sum_{f_j \in N(i)} b_{ij}(\mathbf{n}_i - \mathbf{n}_j) = \mathbf{0}, \quad \forall i, \quad (11)$$

where

$$b_{ij} = \frac{A_i + A_j}{2\nu^2} \phi_\eta(\|\mathbf{c}_i - \mathbf{c}_j\|) \cdot \phi_\mu(\|\mathbf{g}_i - \mathbf{g}_j\|) \cdot \phi_\nu(\|\mathbf{n}_i - \mathbf{n}_j\|). \quad (12)$$

The equations (11) can be solved using fixed-point iteration

$$\mathbf{n}_i^{k+1} = \frac{A_i \hat{\mathbf{n}}_i + \lambda \sum_{f_j \in N(i)} b_{ij}^k \mathbf{n}_j^k}{A_i + \lambda \sum_{f_j \in N(i)} b_{ij}^k}, \quad (13)$$

with $b_{ij}^k = \frac{A_i + A_j}{2\nu^2} \phi_\eta(\|\mathbf{c}_i - \mathbf{c}_j\|) \cdot \phi_\mu(\|\mathbf{g}_i - \mathbf{g}_j\|) \cdot \phi_\nu(\|\mathbf{n}_i^k - \mathbf{n}_j^k\|)$. In this way, the updated normal \mathbf{n}_i^{k+1} of a face f_i is a convex combination of its initial normal $\hat{\mathbf{n}}_i$ and the current normals $\{\mathbf{n}_j^k\}$ of the faces in its neighborhood. The convex combination coefficient for a neighboring face normal \mathbf{n}_j^k depends on both the (static) difference between the guidance normals $(\mathbf{g}_i, \mathbf{g}_j)$ on the two faces, and the (dynamic) difference between their current normals $(\mathbf{n}_i^k, \mathbf{n}_j^k)$, hence the name *static/dynamic* filter. Moreover, there is an interesting connection between the fixed-point iteration and the MM

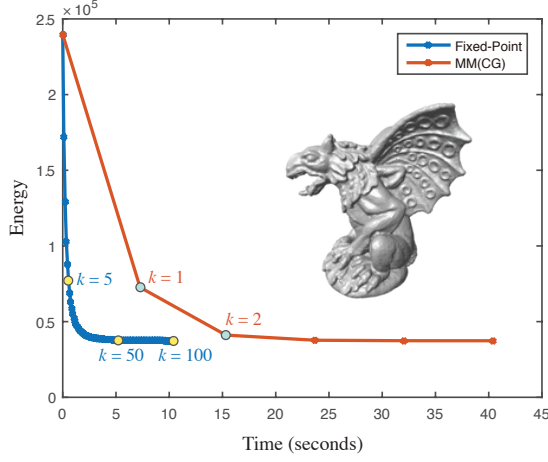


Fig. 3. The change of target energy E_{SD} for the Gargoyle model with respect to the computational time, using the fixed-point iteration solver (13) for 100 iterations and the MM algorithm for 5 iterations, respectively. The fixed-point iteration solver takes much shorter time per iteration, and reduces the energy to a value close to the solution within a fraction of the time for one MM iteration.

algorithm: the iteration (13) is a single step of Jacobi iteration for solving the MM linear system (10).

It can be shown that the fixed-point iteration monotonically decreases the target function value until it converges to a local minimum. A proof is provided in the supplementary material. Moreover, the face normal updates are trivially parallelizable, enabling speedup on GPUs and multi-core CPUs. Our fixed-point solver is significantly faster than the MM algorithm, as shown in Fig. 3 where we compare the change of target energy with respect to the computational time between the two solvers. Here the spatial Gaussian parameter η is set to three times the average distance between neighboring face centroids in the mesh. To achieve the best performance for the MM solver, we tested two strategies for solving the MM linear system: 1) pre-computing symbolic Cholesky factorization of the system matrix, and performing numerical factorization in each iteration, with the three right-hand-sides for x -, and y - and z -coordinates solved in parallel; 2) conjugate gradient method with parallel sparse matrix-vector multiplication. Due to the large neighborhood size, the MM system matrix has a large number of non-zeros in each row, and the Cholesky factorization approach is much more time-consuming than conjugate gradient. Therefore, we only show the conjugate gradient timing. We can see that the fixed-point solver is much more efficient than the MM algorithm, drastically reducing the energy to a value close to the solution within a fraction of the computational time for one MM iteration. This phenomenon is observed in our experiments with other models as well. Detailed results are provided in the supplemental materials.

Enforcing unit normal constraints. The target function E_{SD} in Eq. (4) adapts the SD filter to mesh face normals in a straightforward way, but fails to recognize the requirement that all normals should lie on the unit sphere. In fact, starting from the unit face normals of the input mesh, E_{SD} can be decreased by simply shrinking the normals without changing their directions, and the filtered normals can have different lengths across the mesh. In other words, without

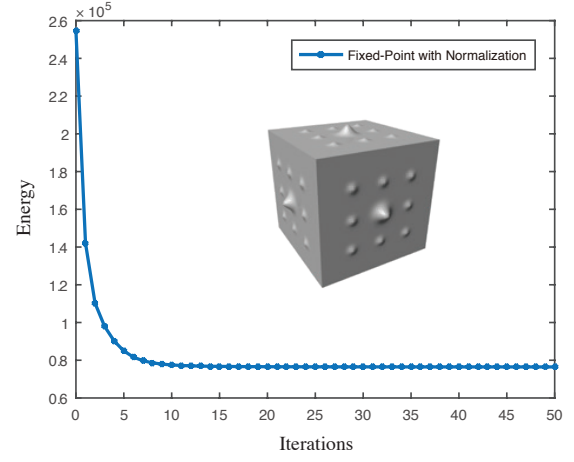


Fig. 4. The change of the modified target energy \bar{E}_{SD} with respect to the number of iterations, using our fixed-point iteration solver (15) for the cube model in Fig. 8. The solver rapidly decreases the energy within a small number of iterations.

the requirement of unit normals, the difference between two normal vectors is not a reliable measure of the deviation between their directions, which makes E_{SD} less effective for controlling the filter. To resolve this issue, we derive a new target function \bar{E}_{SD} for optimization, by substituting each normal vector $\mathbf{n}_i (i = 1, \dots, n_f)$ in E_{SD} with its normalization $\bar{\mathbf{n}}_i = \mathbf{n}_i / \|\mathbf{n}_i\|$. However, such normalization increases the nonlinearity of the problem and makes its numerical solution more challenging. The MM algorithm is no longer applicable, because the quadratic surrogate in Eq. (7) does not hold here. On the other hand, the fixed-point iteration solver can be slightly modified to minimize \bar{E}_{SD} efficiently. At a local minimum, the first-order optimality condition $\partial \bar{E}_{SD} / \partial \mathbf{n}_i = \mathbf{0}$ amounts to

$$(\mathbf{I}_3 - \bar{\mathbf{n}}_i \bar{\mathbf{n}}_i^T) \left(\bar{\mathbf{n}}_i - \frac{A_i \hat{\mathbf{n}}_i + \lambda \sum_{f_j \in N(i)} \bar{b}_{ij} \bar{\mathbf{n}}_j}{A_i + \lambda \sum_{f_j \in N(i)} \bar{b}_{ij}} \right) = \mathbf{0}, \quad (14)$$

where

$$\bar{b}_{ij} = \frac{A_i + A_j}{2\nu^2} \phi_\eta(\|\mathbf{c}_i - \mathbf{c}_j\|) \cdot \phi_\mu(\|\mathbf{g}_i - \mathbf{g}_j\|) \cdot \phi_\nu(\|\bar{\mathbf{n}}_i - \bar{\mathbf{n}}_j\|),$$

and \mathbf{I}_3 is the 3×3 identity matrix. The matrix $\mathbf{I}_3 - \bar{\mathbf{n}}_i \bar{\mathbf{n}}_i^T$ represents the projection onto the subspace orthogonal to $\bar{\mathbf{n}}_i$. Geometrically, condition (14) means that the linear combination $A_i \hat{\mathbf{n}}_i + \lambda \sum_{f_j \in N(i)} \bar{b}_{ij} \bar{\mathbf{n}}_j$ must be parallel to $\bar{\mathbf{n}}_i$. Therefore, we can update \mathbf{n}_i via

$$\mathbf{n}_i^{k+1} = \frac{A_i \hat{\mathbf{n}}_i + \lambda \sum_{f_j \in N(i)} \bar{b}_{ij}^k \bar{\mathbf{n}}_j^k}{\|A_i \hat{\mathbf{n}}_i + \lambda \sum_{f_j \in N(i)} \bar{b}_{ij}^k \bar{\mathbf{n}}_j^k\|}, \quad (15)$$

with $\bar{b}_{ij}^k = \frac{A_i + A_j}{2\nu^2} \phi_\eta(\|\mathbf{c}_i - \mathbf{c}_j\|) \cdot \phi_\mu(\|\mathbf{g}_i - \mathbf{g}_j\|) \cdot \phi_\nu(\|\bar{\mathbf{n}}_i^k - \bar{\mathbf{n}}_j^k\|)$. Compared with the previous iteration format (13), this simply adds a normalization step after each iteration.

Similar to the previous fixed-point iteration format, the new solver with Eq. (15) is embarrassingly parallel, and rapidly decreases the target function within a small number of iterations (see Fig. 4). In the following, all examples of SD normal filtering are processed using this solver.

Vertex update. After filtering the face normals, we need to update the mesh vertices accordingly. Many existing

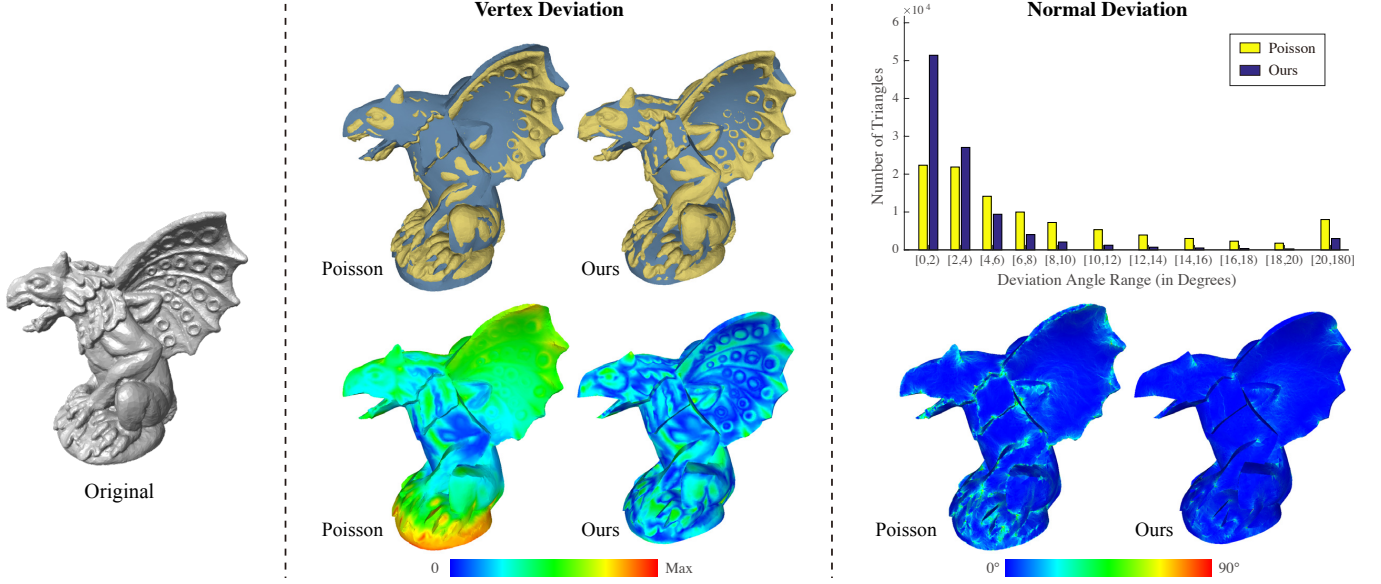


Fig. 5. Comparison between our vertex update method and the Poisson-based update method from [11], based on the same original mesh (left) and target normals. Top center: we align the centroid of each resulting mesh (in blue) with the centroid of the original mesh (in yellow) to minimize the ℓ_2 norm of the deviation between their vertices; the Poisson-based method leads to larger deviation from the original mesh. Bottom center: the deviation between individual vertices after the centroid alignment is visualized via color coding. Right: we evaluate the deviation between the resulting normals and target normals for each face, and visualize its distribution across the mesh via histogram (top right) and color coding (bottom right).

methods do so by enforcing orthogonality between the new edge vectors and the target face normals [31]. Although very efficient, such methods can result in a large number of flipped triangles, because the orthogonality constraint is still satisfied if the updated face normal is opposite to the target one. To address this issue, we propose a new update method that directly enforces the oriented normals as soft constraints, in the same way as [33]. Specifically, for each face f_i with a target oriented unit normal $\hat{\mathbf{n}}_i$, we define \mathcal{C}_i as the feasible set of its vertex positions for which the resulting oriented unit normal is $\hat{\mathbf{n}}_i$. The new vertex positions $\mathbf{V} = [\mathbf{v}_1, \dots, \mathbf{v}_{n_v}]^T \in \mathbb{R}^{n_v \times 3}$ are determined by solving

$$\min_{\mathbf{V}, \mathbf{P}} w \|\mathbf{V} - \mathbf{V}^0\|_F^2 + \sum_{i=1}^{n_f} \|\mathbf{M}\mathbf{V}_{f_i} - \mathbf{P}_i\|_F^2 + \sigma_i(\mathbf{P}_i). \quad (16)$$

Here the first term penalizes the deviation between the new vertex positions and the original vertex positions \mathbf{V}^0 , with $\|\cdot\|_F$ being the Frobenius norm. w is a user-specified positive weight, which is set to 0.001 by default. Matrix $\mathbf{V}_{f_i} \in \mathbb{R}^{3 \times 3}$ stores the vertex positions of face f_i in its rows. Matrix

$$\mathbf{M} = \frac{1}{3} \begin{bmatrix} 2 & -1 & -1 \\ -1 & 2 & -1 \\ -1 & -1 & 2 \end{bmatrix}$$

produces the mean-centered vertex positions for a face. $\mathbf{P}_i \in \mathbb{R}^{3 \times 3}$ are auxiliary variables representing the closest projection of $\mathbf{M}\mathbf{V}_{f_i}$ onto the feasible set \mathcal{C}_i , and σ_i is an indicator function for \mathcal{C}_i , so that

$$\sigma_i(\mathbf{P}_i) = \begin{cases} 0 & \text{if } \mathbf{P}_i \in \mathcal{C}_i, \\ +\infty & \text{otherwise.} \end{cases}$$

The second term of the target function (16) penalizes the violation of oriented normal constraint for each face, using the squared Euclidean distance to the feasible sets. The use of mean-centering matrix \mathbf{M} utilizes the translation-invariance

of the oriented normal constraint, to allow for faster convergence of the solver [33]. Overall, this optimization problem searches for new vertex positions that satisfy the oriented normal constraints as much as possible, while being close the original positions. It is solved via alternating minimization of \mathbf{V} and \mathbf{P} , following the approach of [33]:

- First, we fix \mathbf{V} and update \mathbf{P} . This reduces to a set of separable subproblems, each projecting the current mean-centered vertex positions $\mathbf{M}\mathbf{V}_{f_i}$ of a face to the corresponding feasible set \mathcal{C}_i . Namely, we look for vertex positions \mathbf{P}_i closest to $\mathbf{M}\mathbf{V}_{f_i}$ while achieving the target oriented unit normal $\hat{\mathbf{n}}_i$. The normal condition means that \mathbf{P}_i must lie on a plane orthogonal to $\hat{\mathbf{n}}_i$. Moreover, as the mean of the three vertex positions in $\mathbf{M}\mathbf{V}_{f_i}$ is at the origin, it can be shown that the mean of \mathbf{P}_i must also lie at the origin. As a result, \mathbf{P}_i must lie on a plane that passes through the origin and is orthogonal to $\hat{\mathbf{n}}_i$. The closest projection from $\mathbf{M}\mathbf{V}_{f_i}$ onto this plane can be computed as

$$\bar{\mathbf{P}}_i = \mathbf{M}\mathbf{V}_{f_i}(\mathbf{I}_3 - \hat{\mathbf{n}}_i\hat{\mathbf{n}}_i^T).$$

Let \mathbf{n}_i be the oriented unit normal for the current vertex positions $\mathbf{M}\mathbf{V}_{f_i}$. Then depending on the relation between \mathbf{n}_i and $\hat{\mathbf{n}}_i$, we have two possible solutions for \mathbf{P}_i .

- 1) If $\mathbf{n}_i \cdot \hat{\mathbf{n}}_i \geq 0$, then the oriented unit normal for $\bar{\mathbf{P}}_i$ is $\hat{\mathbf{n}}_i$, and we have $\mathbf{P}_i = \bar{\mathbf{P}}_i$.
- 2) If $\mathbf{n}_i \cdot \hat{\mathbf{n}}_i < 0$, then the oriented unit normal for $\bar{\mathbf{P}}_i$ is $-\hat{\mathbf{n}}_i$. In this case, the solution \mathbf{P}_i degenerates to three colinear points that lie in the plane of $\bar{\mathbf{P}}_i$ and minimizes the distance $\|\mathbf{P}_i - \bar{\mathbf{P}}_i\|_F$. This can be computed as

$$\mathbf{P}_i = \bar{\mathbf{P}}_i(\mathbf{h}\mathbf{h}^T),$$

where \mathbf{h} is the right-singular vector of $\bar{\mathbf{P}}_i$ corresponding to its largest singular value.

The subproblem for each face is independent and can be solved in parallel.

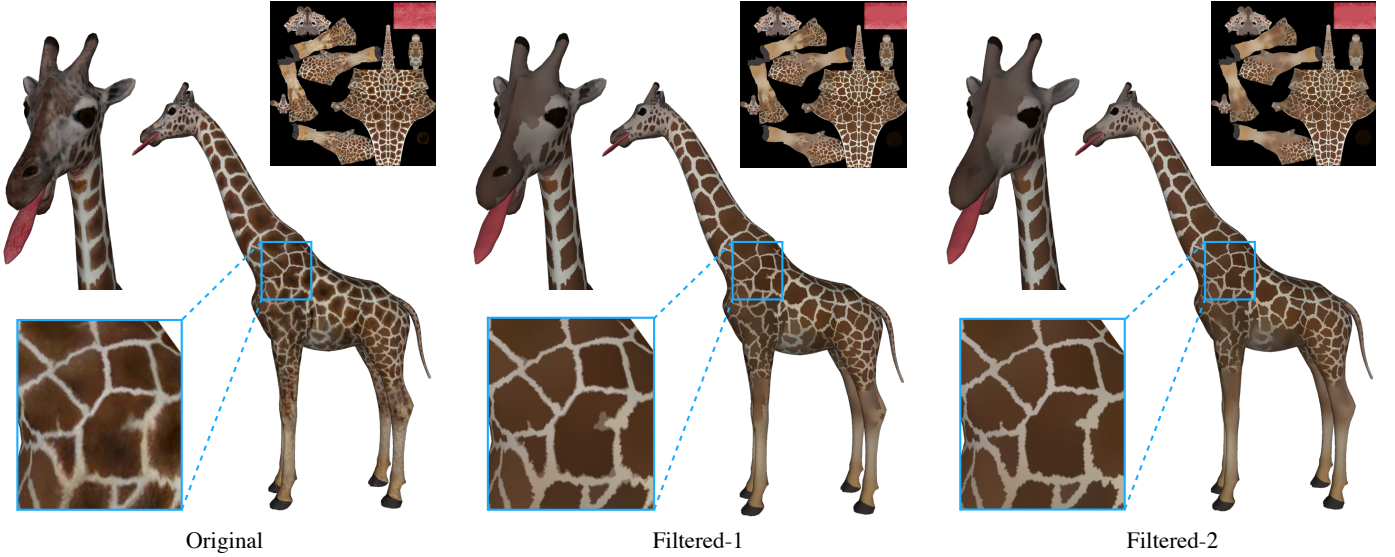


Fig. 6. SD filtering of texture image, which can smooth out features based on their scales on the mesh surface.

- Next, we fix \mathbf{P} and update \mathbf{V} . This is equivalent to

$$\min_{\mathbf{V}} w \|\mathbf{V} - \mathbf{V}^0\|_F^2 + \|\mathbf{K}\mathbf{V} - \mathbf{P}\|_F^2,$$

where the sparse matrix $\mathbf{K} \in \mathbb{R}^{3n_f \times n_v}$ collects the mean-centering matrix coefficients for each face. This amounts to solving a sparse positive definite linear system

$$[w\mathbf{I}_{n_v} + \mathbf{K}^T\mathbf{K}]\mathbf{V} = w\mathbf{V}^0 + \mathbf{K}^T\mathbf{P},$$

where \mathbf{I}_{n_v} is the $n_v \times n_v$ identity matrix. The three right-hand-sides of the system corresponds to the x -, y -, and z -coordinates, and can be solved in parallel. Moreover, the system matrix is fixed during all iterations, thus we can pre-compute its Cholesky factorization to allow for efficient solving in each iteration.

The above alternating minimization is repeated until convergence. We use 20 iterations in all our experiments, which is sufficient to achieve nice results.

Our vertex update method can efficiently compute a new mesh that is consistent with the target face normals, while being close to the original mesh shape. Fig. 5 compares our approach with the vertex update method proposed in [11], which also avoids flipped triangles. Given the target oriented normal for each face, they first rotate the current face to align its oriented normal with the target normal; then the new vertex positions are computed by matching the new face gradients with the rotated ones in a least-squares manner, by solving a Poisson linear system. For each method, we evaluate the deviation between the resulting mesh and the original mesh by aligning their centroids to minimize their ℓ_2 norm of their vertex deviation (shown in the top center of Fig. 5), and then visualizing the deviation of each vertex via color coding. The resulting mesh using our method is noticeably closer to the original mesh, as we explicitly enforce closeness in our target function. This is desirable for many applications such as mesh denoising. In addition, we compute the deviation between the resulting face normals and the target normals, and visualize their distribution using histogram as well as color-coding on the mesh surface. Our

method leads to smaller deviation between the target and the resulting normals. Although the computational time of our method (0.4313 seconds) is higher than the method from [11] (0.1923 seconds), it does not make a significant difference to the total filtering time (see Table 1).

3.2 SD filter for texture colors

Our SD filter can be applied to not only face normals, but also other signals defined on mesh surfaces. One example is RGB colors from texture mapping. Given a texture image and its target triangular mesh, we can use the texture coordinates to identify each pixel i that gets mapped to the surface, as well as its target position $\mathbf{p}_i \in \mathbb{R}^3$ on the mesh. In this way, the texture color becomes a signal defined on the mesh surface, associated with the points $\{\mathbf{p}_i\}$. Let $\{\mathbf{g}_i\}$ be the guidance colors for these points. We compute the filtered texture colors $\{\mathbf{q}_i\}$ by minimizing the target function E_{SD} in Eq. (4), with $\mathbf{c}_i, \mathbf{n}_i$ replaced by $\mathbf{p}_i, \mathbf{q}_i$ respectively, and using the distance between points \mathbf{p}_i to define the neighborhoods and compute the spatial Gaussian weights. Thus the texture colors are filtered according to the metric on the mesh surface instead of the distance on the image plane. The optimization problem is solved using unconstrained fixed-point iteration similar to (13). Fig. 6 shows some examples of texture color filtering.

4 RESULTS AND APPLICATIONS

In this section, we use a series of examples to demonstrate the efficiency and effectiveness of our SD filter, as well as its various applications. We also compare the SD filter with related methods including ℓ_0 optimization [25] and rolling guidance normal filter (RGNF) [11].

Implementation. Our algorithm is implemented in C++, using EIGEN [34] for all linear algebra operations. For filtering of face normals, we run the iterative solver until one of the following conditions is satisfied: 1) the solver reaches the maximum number iterations, which is set to 100 for all our experiments; or 2) the area-weighted




Fig. 7. By repeatedly applying our SD normal filter with different parameters, we can gradually remove the geometry features of increasing scales. Detailed parameter values can be found in the supplemental material.




Fig. 9. Filtering of the Knot model. The SD filter removes the features on each side of the knot and produces smooth appearance, while sharpening the edges between different sides.




Fig. 8. Comparison between the SD normal filter, RGNF, and ℓ_0 optimization, using a cube shape with added features on each side. The colormap shows the deviation between the results and the original cube shape. The result from the SD filter is the closest to the cube.

ℓ_2 norm of normal changes between two consecutive iterations is smaller than a certain threshold angle ϵ , i.e., $\sum_i A_i \|\mathbf{n}_i^{k+1} - \mathbf{n}_i^k\|^2 \leq 4 [\sin(\epsilon/2)]^2 \sum_i A_i$. We set ϵ to 0.2 degrees in all our experiments. For filtering of texture colors, we run the solver for 50 iterations. Unless stated otherwise, all examples are run on a desktop PC with 16GB of RAM and a quad-core CPU at 3.6 GHz, using OpenMP for parallelization.

In all examples, the spatial Gaussian parameter η is specified with respect to the average distance between adjacent face centroids, denoted as l_c . By default, the initial signal is used as the guidance. For more intuitive control of the optimization, we also rescale the user-specified regularizer weight λ according to the value of η . As η increases, the integral of spatial Gaussian ϕ_η on the corresponding




Fig. 10. Scale-aware filtering of the Merlion model, in comparison with ℓ_0 optimization and RGNF. Each row shows comparable results using the three methods, while each column shows the results from one method. The SD filter successfully removes the fine brick lines at the base, while preserving the letters on the base.

face neighborhood also increases, and the relative scale of the regularizer term with respect to the fidelity term grows. Therefore, we compensate for the change of the regularizer scale due to η , by rescaling λ with a factor $(\sum_i A_i) / (\sum_i \sum_{f_j \in N(i)} A_j \phi_\eta(\|\mathbf{c}_i - \mathbf{c}_j\|))$.

The source code for our SD filter is available at <https://github.com/bldeng/MeshSDFilter>. The parameters for the examples in this section can be found in the supplemental material.

Scale-aware and feature-preserving filtering. The SD filter can effectively remove features based on their scales according to the user-specified parameters. This is demonstrated in Fig. 7, where the input models are a sphere




Fig. 11. Coating transfer between two meshes, by first filtering the source mesh to obtain its geometric textures, and applying them to the base cube shape of the target mesh which is also obtained via filtering. The SD filter produces a better result in preserving the sharp edges and recovering the flat faces of the underlying cube shape.

and a cube with additional features of different scales on the surfaces. Using different parameter settings, the SD filter gradually removes the geometry features of increasing scales, while preserving the sharp edges on the cube. Similar scale-aware and feature-preserving effects are observed for filtering of texture colors, as shown in Fig. 6.

In Figs. 8, 9, and 10, we show more examples of scale-aware and feature-preserving filtering of mesh geometry using the SD filter, and compare the results with ℓ_0 optimization method from [25] and RGNF from [11]. We tune the parameters of each method to achieve the best results, while ensuring comparable effects from different methods. In all examples, the SD filter achieves better or similar results compared with RGNF, and outperforms ℓ_0 optimization. In Fig. 8, the input model is a cube with additional features on each face, and the result with the SD filter is the closest to the cube shape. The result with RGNF has larger deviation from the cube shape, because the filtered signals are computed as a combination of the original signals within a neighborhood; as a consequence, when there is large deviation between the input signals and the desired output within a certain region, RGNF may not produce a desirable result inside the region. By default, Wang et al. [11] run RGNF for 5 iterations. To verify its convergence, we run 50 iterations of RGNF and compare the results. We can see that the resulting mesh with 50 RGNF iterations is slightly closer to the cube shape, but the difference is not significant, and the undesirable deviation around the edges remains. For comparison, we also run the SD filter for 50 and 500 iterations respectively. Similar to RGNF, running more iterations of SD filters makes the result slightly closer to the cube shape, but without significant difference. In Fig. 9, the SD filter is able to smooth out the star-shaped features on the knot surface, while enhancing the sharp feature lines between different sides of the knot. Although ℓ_0 optimization also enhances the feature lines, it leads to piecewise flat shapes because the ℓ_0 norm promotes piecewise constant signals. Similar to Fig. 8, we run RGNF for 5 and 50 iterations, and the SD filter for 50 and 500 filters, to compare their results and convergence. For both filters, running for more iterations slightly improve the result, but without significant difference. In Fig. 10, the three methods




Fig. 12. A larger value of λ or η leads to a smoother filtering result. Here each rows shows the filtering results with increasing values of λ or η , while keeping the other parameters fixed.

produce similar results on the Merlion model, while the scale-awareness of the SD filter enables it to remove the fine brick lines at the base while clearly retaining the letters.

Coating transfer. Similar to [11], our filter can be utilized to perform coating transfer [35], which transfers geometric textures from a source mesh to a target mesh. An example is shown in Fig. 11. Here following [11], we first filter the source mesh to remove the letters on its surface, and encode them using the difference between Laplacian coordinates of the mesh before and after the filtering. We then transfer these letters onto the base cube shape of the target mesh which is also obtained by filtering: the encoded letter shapes are added to the Laplacian coordinates of the base cube mesh, and a new mesh is reconstructed from the resulting Laplacian coordinates. For comparison, we also show the results using [11], where the meshes are processed using the RGNF instead of the SD filter. While both methods can extract the source geometric textures and prepare the target cube mesh, the SD filter produces better a result in preserving the sharp edges and recovering the flat faces of the cube shape.

Choice of parameters. Our filter is influenced by four parameters: the regularizer weight λ , the spatial Gaussian parameter η which also influences the neighborhood size, the guidance Gaussian parameter μ , and the range Gaussian parameter ν . For both λ and η , a larger value leads to a smoother result, as shown in Fig. 12. Parameters μ and ν determine which face normals within a neighborhood affects the central face in a fixed-point iteration. For more intuitive control, we propose a method to interactively set the μ and ν parameters. Firstly, the user selects two smooth regions on different sides of a feature intended to be kept. We denote the two regions as P_1 and P_2 , and compute the mean $\tilde{\mathbf{n}}_l$ and variance σ_l ($l = 1, 2$) of the normals within each region via

$$\tilde{\mathbf{n}}_l = \frac{\sum_{i \in P_l} A_i \mathbf{n}_i}{\|\sum_{i \in P_l} A_i \mathbf{n}_i\|}, \quad \sigma_l^2 = \frac{\sum_{i \in P_l} A_i \|\mathbf{n}_i - \tilde{\mathbf{n}}_l\|^2}{\sum_{i \in P_l} A_i}.$$

Then the range of ν is determined using the following strategy: ν should be small enough such that the two mean normals have negligible influence on each other according to the range Gaussian; at the same time, within each region the normals should influence each other such that sharp features

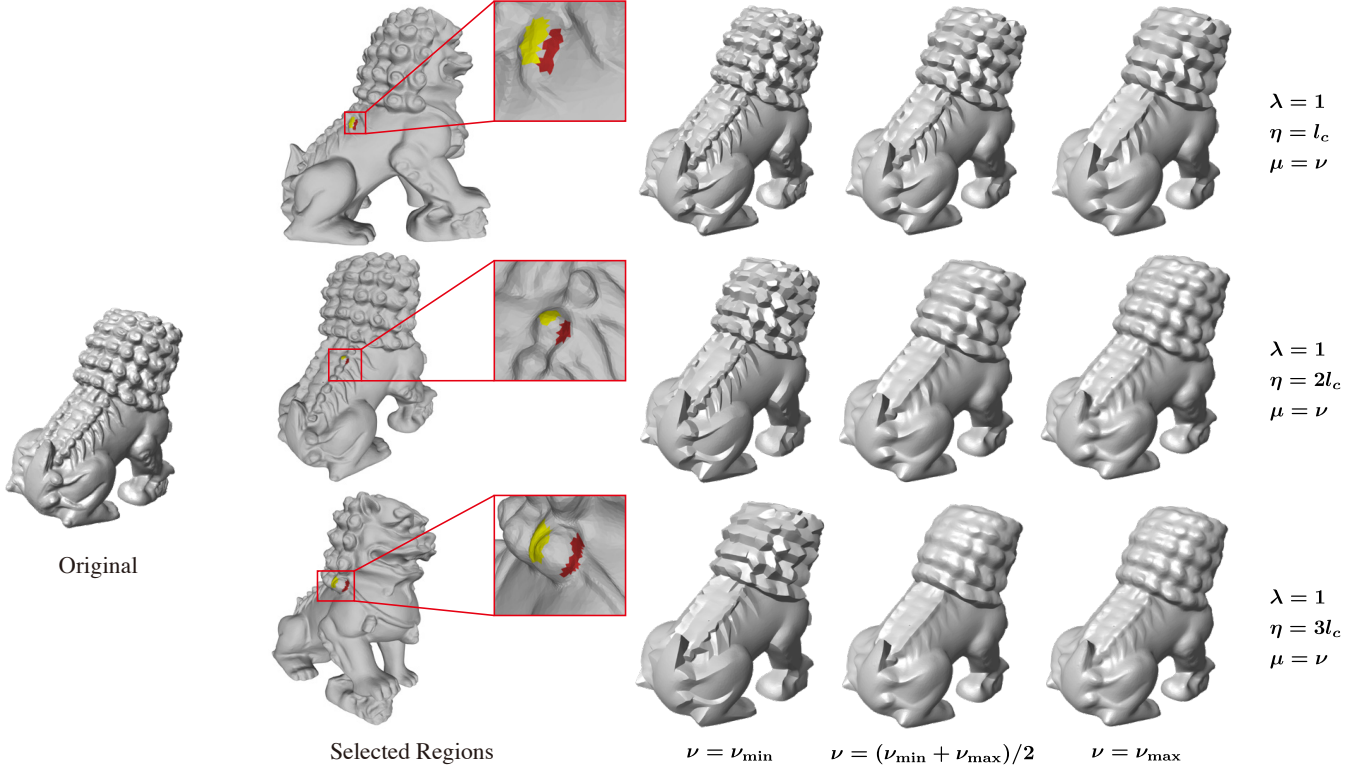


Fig. 13. According to two user-selected regions (shown in red and yellow), we determine upper bound ν_{\max} and ν_{\min} for the ν parameter, which inhibits influence between the face normals from the two regions, and ensures the smoothing of normals within each region. Within this range, a larger ν leads to smoother results, while a smaller ν promotes sharp features between the two selected regions. Here each row shows one pair of selected regions, and the resulting meshes using different values of ν within the corresponding range.

do not emerge. Based on this strategy, we first determine the lower- and upper-bounds of ν via

$$\nu_{\min} = \frac{1}{2} \max(\sigma_1, \sigma_2), \quad \nu_{\max} = \frac{1}{3} \|\tilde{\mathbf{n}}_1 - \tilde{\mathbf{n}}_2\|.$$

If $\nu_{\max} < \nu_{\min}$, then the user needs to select another pair of regions; otherwise, a value between ν_{\min} and ν_{\max} as the parameter ν . In our experiments, good results can often be achieved by choosing $\nu = (\nu_{\min} + \nu_{\max})/2$ and setting $\mu = a \cdot \nu$ with $a \in [1, 10]$. Fig 13 shows the effects of different ν values on the Chinese lion model.

Feature manipulation and enhancement. The scale-awareness of our filter enables us to manipulate mesh details according to their scales. Given an input mesh M , we repeatedly apply the SD normal filter with different parameters to obtain a series of filtered meshes M^{m-1}, \dots, M^1, M^0 , where $M^k (k = 0, \dots, m-2)$ is obtained by filtering M^{k+1} to remove more details. If we denote the input mesh as M^m , then M^0, \dots, M^m forms a coarse-to-fine sequence of meshes, with M^0 being the *base mesh*, and M^m being the original mesh. We encode the difference between two consecutive meshes M^k, M^{k-1} by comparing their corresponding vertex positions and face normals, represented as $\mathbf{d}_i^k = \mathbf{v}_i^{(k)} - \mathbf{v}_i^{(k-1)} (i = 1, \dots, n_v)$ and $\mathbf{f}_j^k = \mathbf{n}_j^{(k)} - \mathbf{n}_j^{(k-1)} (i = 1, \dots, n_f)$, where $\{\mathbf{v}_i^{(k)}\}$ and $\{\mathbf{n}_j^{(k)}\}$ are the vertex positions and face normals of mesh M^k . These differences represent the required deformation for M^{k-1} to introduce the additional details in M^k . They can be linearly combined according to coefficients $\alpha = [\alpha_1, \dots, \alpha_m]$

and added to the face normals and vertex positions of the base mesh, to derive the target vertex positions $\{\hat{\mathbf{v}}_i\}$ and target face normals $\{\hat{\mathbf{n}}_j\}$ for a new mesh \bar{M} :

$$\hat{\mathbf{v}}_i = \mathbf{v}_i^{(0)} + \sum_{k=1}^m \alpha_k \mathbf{d}_i^k, \quad \hat{\mathbf{n}}_j = \frac{\mathbf{n}_j^{(0)} + \sum_{k=1}^m \alpha_k \mathbf{f}_j^k}{\|\mathbf{n}_j^{(0)} + \sum_{k=1}^m \alpha_k \mathbf{f}_j^k\|}. \quad (17)$$

Note that the target vertex positions and target face normals are often incompatible. To combine the two conditions, we determine the new mesh by solving the same optimization problem as our vertex update (16), with the matrix \mathbf{V}^0 in the target function storing the target vertex positions. In this way, the linear combination coefficients α indicate the contribution of geometric features from the original model within a certain range of scales. By changing the value of α , a user can control geometric features according to their scales. Setting all coefficients to 1 recovers the original mesh, while setting a coefficient to a value different from 1 can boost or attenuate the features of the corresponding scales. Moreover, the linear system matrix for the optimization problem is fixed regardless of the value of α , and only needs to be pre-factorized once. Afterwards, the user can choose any value of α , and the resulting mesh can be efficiently computed using the pre-factorized system. This allows the user to interactively explore different linear combination coefficients to achieve desirable results. Figs. 1, 14, and 15 show examples of new meshes created in this manner. We can see that the coarse-to-fine sequence of meshes captures the geometrical features of different scales, which are effectively manipulated using the linear combination coefficients. In




Fig. 14. Geometry feature manipulation and enhancement for the Armadillo model, by controlling the contribution from features of different scales. Left: a coarse-to-fine sequence of meshes, obtained by repeatedly applying the SD normal filter with different parameters. Right: new meshes generated using linearly combined target vertex positions and target normals (Eq. (17)), with the combination coefficients shown below each result.




Fig. 15. Geometry feature manipulation and enhancement for the Welsh Dragon model. Left: the coarse-to-fine sequence of meshes resulting from SD normal filtering. Right: generated new meshes and their corresponding linear combination coefficients.

some application scenarios, it is desirable to only modify the features within a certain region on the surface. In this case, the target vertex positions and target normals are linearly combined only within user-selected regions; outside these regions they remain the same as the original mesh. Fig. 17 shows such an example, where we enhance the features of the Gargoyle model in a user-selected local region. Another example is shown in Fig. 18, where a 3D human face model is locally enhanced.

Similarly, we can manipulate and enhance texture colors using the SD filter, as shown in Fig. 16. We first filter the input texture colors T incrementally to derive a coarse-to-fine

sequence $T^0, \dots, T^{m-1}, T^m = T$. Then new texture colors are computed via linear combination with coefficients α :

$$T_{\text{new}} = T^0 + \sum_{k=1}^m \alpha_k (T^k - T^{k-1}). \quad (18)$$

Out-of-bound colors in T_{new} are reset to their closest valid values.

Mesh denoising. By constructing appropriate guidance signals, our SD filter can also be applied for mesh denoising, as shown in Fig. 19. Here we repeatedly apply the SD normal filter to each input mesh to remove the noise. In each run of the filter, the guidance normals are computed from the cur-



Fig. 16. Texture image feature enhancement using the SD filter.

rent mesh using the patch-based approach proposed in [10]. The results are evaluated using the average normal deviation δ and the average vertex deviation $\mathcal{D}_{\text{mean}}$ from the ground-truth mesh as proposed in [10]. In addition, similar to [9], we measure the perceptual difference between the denoised mesh and the ground truth using the spatial error term of the STED distance proposed in [36], computed with one-ring vertex neighborhood. Fig. 19 compares our denoising results with the guided mesh normal filtering (GMNF) method from [10]. The results from the two methods are quite close, with similar error metric values. Detailed parameter settings are provided in the supplementary materials.

Performance. Using parallelization, our SD filter can compute the results efficiently. Table 1 provides the representative computation time of the SD normal filter on different models, showing the timing for each part of the algorithm:

- T_1 : the pre-processing time for finding the neighborhood;
- T_2 : the average timing of one iteration;
- T_3 : the timing for mesh vertex update;
- T_{total} : The total timing for the whole filtering process.

For meshes with less than 100K faces and with $\eta \leq 3l_c$, the whole process typically takes only a few seconds.




Fig. 17. Detail enhancement of the Gargoyle model in a local region selected by the user (shown in yellow), using the same coarse-to-fine mesh sequence as Fig. 1.

5 DISCUSSION AND CONCLUSION

We present the SD filter for triangular meshes, which is formulated as an optimization problem with a target energy that combines a quadratic fidelity term and a nonconvex robust regularizer. We develop an efficient fixed-point iteration solver for the problem, enabling the filter to be applied for interactive applications. Our SD filter generalizes the joint bilateral filter, combining the static guidance with a dynamic guidance that is derived from the current signal values. Thanks to the joint static/dynamic guidance, the SD filter is robust, feature-preserving and scale-aware, producing state-of-the-art results for various geometry processing problems.

Although our solver can incorporate simple constraints such as unit length for normal vectors, we do not consider




Fig. 18. Local feature enhancement on a human face model, with the local region and the combination coefficients annotated below each result.

TABLE 1
Computational Time (in Seconds) for Different Parts of the SD Normal Filtering Method

Model	#Faces	λ	η	μ	ν	T_1	T_2	T_3	T_{total}
Armadillo	43K	2	$2.5l_c$	1.5	0.45	0.57	0.036	0.19	1.45
Cube	49K	1×10^6	$5l_c$	2.5	0.4	2.03	0.15	0.37	3.42
Sphere	60K	100	$3l_c$	1.5	0.17	1.10	0.065	0.34	2.88
Duck	68K	10	$2l_c$	2.5	0.3	0.61	0.031	0.33	1.69
Knot	100K	10	$2l_c$	2.5	0.8	1.30	0.07	0.50	5.69
Gargoyle	100K	5	$3l_c$	10	0.42	1.90	0.12	0.35	8.17
Merlion	566K	10	$2.5l_c$	2	0.26	6.96	0.49	5.24	32.22
Welsh Dragon	2.21M	10	$1.5l_c$	20	0.35	19.85	0.49	36.17	71.26

global conditions for the signals. For example, we do not ensure the integrability of normals, i.e., the existence of a mesh whose face normals match the filter results; as a result, some parts of the updated mesh may not be consistent with the filtered normals. Neither do we consider the prevention of self intersection of the updated mesh. Due to the local nature of our fixed-point iteration, it is not easy to incorporate such global constraints into the solver. A possible remedy is to introduce a separate step to enforce these conditions after a few iterations. A more in-depth investigation into such global conditions will be an interesting problem.

Depending on the parameters, our filter can sometimes create sharp features in regions that are originally smooth. For example, this can be seen in the top right image of Fig. 10, around the scales of the Merion model. One potential solution is to use spatially varying filter parameters that adapt to local feature sizes and user preferences, which we will leave as future work.

Although we only consider filtering of face normals and texture colors in this paper, our formulation is general enough to be applied for other scenarios. In the future, we would like to extend the filter to other geometry signals such as curvatures and shape operators, and to other geometric representations such as point clouds and implicit surfaces. Also worth investigating is the application of fixed-point iteration to similar image filtering problems, especially where direct linear solve is too slow due to the problem scale or the neighborhood size.

ACKNOWLEDGMENTS

We thank Yang Liu for providing the implementation of RGNF. The Welsh Dragon mesh model was released by Bangor University, UK, for Eurographics 2011. This work was supported by the National Key R&D Program of China (No. 2016YFC0800501), the National Natural Science Foundation of China (No. 61672481, No. 61672482 and No. 11626253), and the One Hundred Talent Project of the Chinese Academy of Sciences.

REFERENCES

- [1] C. Tomasi and R. Manduchi, "Bilateral filtering for gray and color images," ser. ICCV '98, 1998.
- [2] E. Eisemann and F. Durand, "Flash photography enhancement via intrinsic relighting," *ACM Trans. Graph.*, vol. 23, no. 3, pp. 673–678, 2004.
- [3] G. Petschnigg, R. Szeliski, M. Agrawala, M. Cohen, H. Hoppe, and K. Toyama, "Digital photography with flash and no-flash image pairs," *ACM Trans. Graph.*, vol. 23, no. 3, pp. 664–672, 2004.
- [4] H. Cho, H. Lee, H. Kang, and S. Lee, "Bilateral texture filtering," *ACM Trans. Graph.*, vol. 33, no. 4, pp. 128:1–128:8, 2014.
- [5] Q. Zhang, X. Shen, L. Xu, and J. Jia, "Rolling guidance filter," in *Computer Vision—ECCV 2014*. Springer, 2014, pp. 815–830.
- [6] S. Fleishman, I. Drori, and D. Cohen-Or, "Bilateral mesh denoising," *ACM Trans. Graph.*, vol. 22, no. 3, 2003.
- [7] T. R. Jones, F. Durand, and M. Desbrun, "Non-iterative, feature-preserving mesh smoothing," *ACM Trans. Graph.*, vol. 22, no. 3, pp. 943–949, 2003.
- [8] Y. Zheng, H. Fu, O.-C. Au, and C.-L. Tai, "Bilateral normal filtering for mesh denoising," *IEEE Trans. Vis. Comput. Graphics*, vol. 17, no. 10, pp. 1521–1530, 2011.
- [9] J. Solomon, K. Crane, A. Butscher, and C. Wojtan, "A general framework for bilateral and mean shift filtering," *arXiv preprint arXiv:1405.4734*, 2014.





Fig. 19. Comparison between denoising results using the SD normal filter, and the guided mesh normal filtering (GMNF) method [10]. The results from the two methods are close, with similar error metric values.


- [10] W. Zhang, B. Deng, J. Zhang, S. Bouaziz, and L. Liu, "Guided mesh normal filtering," *Comput. Graph. Forum*, vol. 34, no. 7, pp. 23–34, 2015.
- [11] P. Wang, X. Fu, Y. Liu, X. Tong, S. Liu, and B. Guo, "Rolling guidance normal filter for geometric processing," *ACM Trans. Graph.*, vol. 34, no. 6, p. 173, 2015.
- [12] B. Ham, M. Cho, and J. Ponce, "Robust image filtering using joint static and dynamic guidance," in *CVPR*, 2015.
- [13] G. Taubin, "A signal processing approach to fair surface design," ser. SIGGRAPH '95, 1995, pp. 351–358.
- [14] M. Desbrun, M. Meyer, P. Schröder, and A. H. Barr, "Implicit fairing of irregular meshes using diffusion and curvature flow," ser. SIGGRAPH '99, 1999, pp. 317–324.
- [15] G. Taubin, "Linear anisotropic mesh filtering," IBM Research, IBM Research Report RC22213(W0110-051), 2001.
- [16] Y. Ohtake, A. Belyaev, and I. Bogaevski, "Mesh regularization and adaptive smoothing," *Computer-Aided Design*, vol. 33, no. 11, pp. 789–800, 2001.
- [17] M. Chuang and M. M. Kazhdan, "Interactive and anisotropic geometry processing using the screened poisson equation," *ACM Trans. Graph.*, vol. 30, no. 4, pp. 57:1–57:10, 2011.
- [18] S. Paris, P. Kornprobst, J. Tumblin, and F. Durand, *Bilateral filtering: Theory and applications*. Now Publishers Inc, 2009.
- [19] T. Jones, F. Durand, and M. Zwicker, "Normal improvement for point rendering," *IEEE Computer Graphics and Applications*, vol. 24, no. 4, pp. 53–56, 2004.
- [20] C. C. Wang, "Bilateral recovering of sharp edges on feature-insensitive sampled meshes," *IEEE Trans. Vis. Comput. Graphics*, vol. 12, no. 4, pp. 629–639, 2006.
- [21] J. Kopf, M. F. Cohen, D. Lischinski, and M. Uyttendaele, "Joint bilateral upsampling," *ACM Trans. Graph.*, vol. 26, no. 3, p. 96, 2007.
- [22] L. Xu, C. Lu, Y. Xu, and J. Jia, "Image smoothing via L_0 gradient minimization," *ACM Trans. Graph.*, vol. 30, no. 6, pp. 174:1–174:12, 2011.
- [23] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Phys. D*, vol. 60, no. 1-4, pp. 259–268, 1992.
- [24] G. Taubin, "Introduction to geometric processing through optimization," *IEEE Computer Graphics and Applications*, vol. 32, no. 4, pp. 88–94, 2012.
- [25] L. He and S. Schaefer, "Mesh denoising via L_0 minimization," *ACM Trans. Graph.*, vol. 32, no. 4, pp. 64:1–64:8, 2013.
- [26] H. Zhang, C. Wu, J. Zhang, and J. Deng, "Variational mesh denoising using total variation and piecewise constant function space," *IEEE Transactions on Visualization and Computer Graphics*, vol. 21, no. 7, pp. 873–886, 2015.
- [27] P. Perona and J. Malik, "Scale-space and edge detection using anisotropic diffusion," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 12, no. 7, pp. 629–639, 1990.
- [28] J. Sun, M. Ovsjanikov, and L. Guibas, "A concise and provably informative multi-scale signature based on heat diffusion," in *Proceedings of the Symposium on Geometry Processing*, 2009, pp. 1383–1392.
- [29] B. Vallet and B. Lévy, "Spectral geometry processing with manifold harmonics," *Computer Graphics Forum*, vol. 27, no. 2, pp. 251–260, 2008.
- [30] H. Zhang, O. Van Kaick, and R. Dyer, "Spectral mesh processing," *Computer Graphics Forum*, vol. 29, no. 6, pp. 1865–1894, 2010.
- [31] X. Sun, P. L. Rosin, R. R. Martin, and F. C. Langbein, "Fast and effective feature-preserving mesh denoising," *IEEE Trans. Vis. Comput. Graphics*, vol. 13, no. 5, pp. 925–938, 2007.
- [32] F. Pukelsheim, "The three sigma rule," *The American Statistician*, vol. 48, no. 2, pp. 88–91, 1994.
- [33] S. Bouaziz, M. Deuss, Y. Schwartzburg, T. Weise, and M. Pauly, "Shape-up: Shaping discrete geometry with projections," *Computer Graphics Forum*, vol. 31, no. 5, pp. 1657–1667, 2012.
- [34] G. Guennebaud, B. Jacob et al., "Eigen v3," <http://eigen.tuxfamily.org>, 2010.
- [35] O. Sorkine, D. Cohen-Or, Y. Lipman, M. Alexa, C. Rössl, and H.-P. Seidel, "Laplacian surface editing," in *Proceedings of the 2004 Eurographics/ACM SIGGRAPH Symposium on Geometry Processing*, ser. SGP '04, 2004, pp. 175–184.
- [36] L. Vasa and V. Skala, "A perception correlated comparison method for dynamic meshes," *IEEE Transactions on Visualization and Computer Graphics*, vol. 17, no. 2, pp. 220–230, 2011.



Juyong Zhang is an associate professor in the School of Mathematical Sciences at University of Science and Technology of China. He received the BS degree from University of Science and Technology of China in 2006, and the PhD degree from Nanyang Technological University, Singapore. His research interests include computer graphics, computer vision, and numerical optimization. He is a member of the IEEE.



Bailin Deng is a lecturer in the School of Computer Science and Informatics at Cardiff University. He received the BEng degree in computer software (2005) and the MSc degree in computer science (2008) from Tsinghua University (China), and the PhD degree in technical mathematics from Vienna University of Technology (Austria). His research interests include geometry processing, numerical optimization, computational design, and digital fabrication. He is a member of the IEEE.



Yang Hong is currently working toward a master's degree in the School of Mathematical Sciences, University of Science and Technology of China. His research interests include geometric & image processing and 3D modeling.



Yue Peng received her BS degree from the University of Science and Technology of China in 2016. Currently, she is working toward a master's degree in the USTC. Her research interests include computer graphics and geometry processing.



Wenjie Qin received the bachelor's degree from the University of Science and Technology of China in 2012. He is currently working toward a master degree in the University of Science and Technology of China. His research interests include geometric processing and filtering.



Ligang Liu is a Professor at the School of Mathematical Sciences, University of Science and Technology of China. His research interests include digital geometric processing, computer graphics, and image processing. He serves as the associated editors for journals of IEEE Transactions on Visualization and Computer Graphics, IEEE Computer Graphics and Applications, Computer Graphics Forum, Computer Aided Geometric Design, and The Visual Computer. He served as the conference co-chair of GMP 2017 and the

program co-chairs of GMP 2018, CAD/Graphics 2017, CVM 2016, SGP 2015, and SPM 2014. His research works could be found at his research website: <http://staff.ustc.edu.cn/lgliu>.