DOI: 10.1111/cgf.12625
Eurographics Conference on Visualization (EuroVis) 2015
H. Carr, K.-L. Ma, and G. Santucci
(Guest Editors)

Volume 34 (2015), Number 3

Edge-Aware Volume Smoothing Using L_0 Gradient Minimization

Qichao Wang, Yubo Tao†, Hai Lin

State Key Laboratory of CAD&CG, Zhejiang University, P.R. China

Abstract

In volume visualization, noise in regions of homogeneous material and at boundaries between different materials poses a great challenge in extracting, analyzing and rendering features of interest. In this paper, we present a novel volume denoising / smoothing method based on the L_0 gradient minimization framework. This framework globally controls how many voxels with a non-zero gradient are in the result in order to approximate important features' structures in a sparse way. This procedure can be solved quickly by the alternating optimization strategy with half-quadratic splitting. While the proposed L_0 volume gradient minimization method can effectively remove noise in homogeneous materials, a blurring-sharpening strategy is proposed to diminish noise or smooth local details on the boundaries. This generates salient features with smooth boundaries and visually pleasing structures. We compare our method with the bilateral filter and anisotropic diffusion, and demonstrate the effectiveness and efficiency of our method with several volumes in different modalities.

Categories and Subject Descriptors (according to ACM CCS): I.4.10 [Image Processing and Computer Vision]: Image Representation—Volumetric

1. Introduction

Volume visualization has been widely used for feature analysis and presentation in different fields, ranging from medicine and biology to engineering and geophysics. These volumes are usually generated through scanners in a diversity of modalities, such as Computed Tomography (CT), Magnetic Resonance Imaging (MRI) and Ultrasound. However, even with high-resolution scanners, aliasing and noise are inevitably introduced during the volume acquisition procedure. This noise could corrupt salient features in terms of intensity, especially low-contrast structures and small objects, and this makes intensity-based volume analysis methods, such as transfer functions and iso-surface extraction, fail to produce accurate feature information and yield visual confusion and misunderstanding in the rendered result. Therefore, noise is a great challenge in volume visualization, and denoising / smoothing is often required in the pre-processing step to reduce noise while preserving important features.

There have been many attempts to deal with the volume smoothing problem. Filtering is one of these frameworks for volume smoothing. The Gaussian filter is widely used in volume smoothing, but boundaries are also blurred. The bilateral filter [TM98] considers both the spatial distance and intensity difference between two pixels to better preserve edges. Another popular framework in volume smoothing is anisotropic diffusion, originally introduced by Perona and Malik [PM90]. Anisotropic diffusion iteratively performs edge-preserving local filtering to suppress noise, and the original scheme tends to over-sharpen edges. Recently, a new edge-aware anisotropic diffusion model [HM10] defines the boundaries based on the directional second derivative along the gradient to denoise the volume while preserving boundaries and fine tubular structures. However, anisotropic diffusion is a relatively slowly converging non-linear iterative process. As previous volume smoothing methods are mostly based on local filters, they result in more or less boundaries blurriness and are hard to globally maintain and possibly sharpen the most prominent boundaries while enabling a reliable primary-structure representation.

Most recently, a novel L_0 minimization framework was introduced to globally smooth images [XLXJ11]. The L_0 nor-

† Corresponding Author: Yubo Tao (taoyubo@cad.zju.edu.cn)

m of a vector is the number of non-zero entities. It directly measures the information sparsity and is very powerful to abstract the main structural information. Compared with previous edge-preserving methods, the resulted signal suppresses noise / details, and retains and sharpens significant edges. At the same time, this method does not introduce any edge blurriness as it is not based on local filters.

This paper presents a new volume denoising / smoothing method based on the L_0 gradient minimization framework. As the gradient magnitude is highly related to structural information, especially in medical datasets [KD98], we can restrict the number of intensity changes among neighboring voxels in the context of L_0 gradient minimization. This leads to an unconventional global optimization procedure based on the discrete metric, which counts the number of voxels with a non-zero gradient. We employ the alternating optimization strategy with half-quadratic splitting to optimize this procedure. The proposed L_0 volume gradient minimization is greatly helpful for suppressing less important details and noise with low gradient magnitudes, and sharpening and enhancing salient features at the same time.

The L_0 volume gradient minimization can perfectly smooth details and noise in relatively homogenous materials, but it cannot effectively suppress details and noise on the boundaries due to their relatively high gradient magnitudes. This results in visual artifacts and misunderstanding during feature analysis. As the proposed L_0 volume gradient minimization can sharpen boundaries, this paper proposes a blurring-sharpening strategy to achieve feature boundary denoising / smoothing. A linear filter, such as the Gaussian filter, is firstly used to smooth noise on the boundaries, and then the proposed L_0 volume gradient minimization is applied to the smoothed volume to diminish the blurred noise and sharpen the blurred boundaries. In this way, the smoothed result contains salient features with smooth boundaries and visually pleasing structures, and it is also more consistent with the original volume. To the best of our knowledge, our method is the first attempt to apply L_0 gradient minimization in volume denoising / smoothing.

The paper is structured as follows. Related work is discussed in Section 2, while Section 3 describes the proposed volume smoothing method. In Section 4, we present the blurring-sharpening strategy to deal with noise at boundaries, while Section 5 presents and discusses the smoothed / denoised results of several volume datasets in different modalities. Finally, we draw conclusions in Section 6.

2. Related Work

Smoothing is a widely investigated research area in many fields, such as image and mesh processing, and volume visualization. This section briefly reviews three main categories of smoothing: local filtering, diffusion, and non-local methods.

Local filtering usually replaces the value with a weighted mean of its neighbors, and the weights are often related to the spatial distance and / or the value difference [MMMY97, MMK*98, SH05, MAD14]. The Gaussian filter is effective at removing noise from the volume, but is not able to preserve boundaries of features, resulting in an over-smoothing effect around the boundary. Many edge-aware local filtering methods have been proposed to overcome this problem. One popular method is bilateral filtering [TM98], which performs a weighted averaging of the values in a window based on both space and range distances for image smoothing. This method has been extended to mesh smoothing [FDCO03]. The bilateral filter can be directly extended to a 3D volume to filter out small changes in intensity and preserve strong boundaries, and it comes at the expense of tradeoff between details suppression and edge preservation [FFLS08].

Diffusion is an iterative process based on local filtering. Anisotropic diffusion is often used for suppressing noise while preserving important structures. Its capability depends on the stopping function to prevent smoothing from crossing strong edges, otherwise this will cause edge blurriness. It was originally introduced by Perona and Malik [PM90], and can be used for intense smoothing while preserving prominent edges between the smooth regions. Krissian et al. [KMA97] presented an anisotropic diffusion model to perform diffusion along the direction of the gradient and the minimum and maximum curvature. Krissian and Aja-Fernández [KAF09] introduced a noise-driven anisotropic diffusion that combines local linear minimum mean square error filters and partial differential equations to remove Rician noise from a 3D MRI volume. Recently, a novel edge aware anisotropic diffusion model proposed by Hossain and Möller [HM10] uses the directional second derivative to define material boundaries, and thus it can preserve material boundaries as well as fine tubular structures. As these anisotropic diffusion methods usually apply local filtering in each iteration, they blur edges to some extent in the final result. The proposed method avoids local filtering and is a global approach to smooth the low gradient magnitude regions while retaining or even sharpening salient boundaries.

Non-local methods attempt to remove noise without blurring salient edges [BCM05]. Topological denoising filters [WGS10, JWS12] first globally extract and filter the extrema of a 2D scalar field, and then obtain the denoised result through the solution of a discrete optimization problem. Günther et al. [GJR*14] extended the topology-controlled denoising method to globally preserve the minima and maxima of large 2D data and medium-sized 3D data by introducing a novel domain decomposition approach. However, the computation is very time consuming, several hours for a 128^3 volume. On the other hand, L_0 minimization is a new well-defined and powerful smoothing framework based on the principle of discretely counting spatial intensity changes. Compared with local smoothing methods, it can diminish low gradient amplitude details and globally preserve and en-

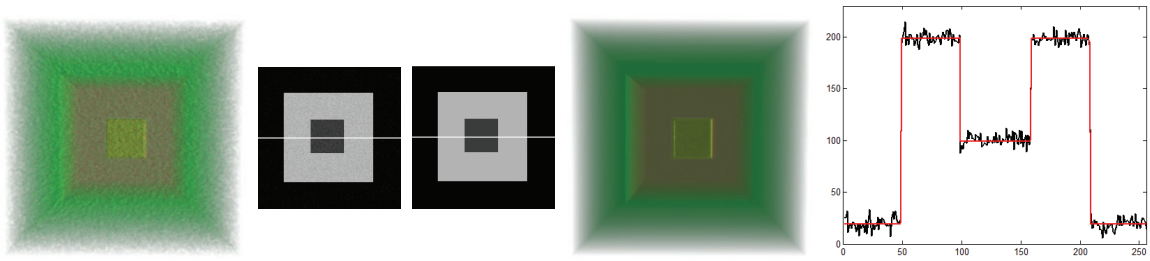

(a) (b) (c)

Figure 1: *A denoising example of the L_0 volume gradient minimization method. (a) The noisy volume, corrupted by synthetic noise. (b) The smoothed volume after our method. (c) The profile plot is a 1D scanline plot of the scalar values taken from a slice as indicated by the white line in the 2D images. The black curve is a plot of the noisy scalar values, while the red one is the plot of the smoothed scalar values. Our method smoothes noise while preserving the most significant sharp edges.*

hance salient features, even if the data is low-contrast or important features are small. Xu et al. [XLXJ11] utilized the L_0 minimization framework to image smoothing to progressively suppress details. However, smoothness of feature boundaries is not taken into account. He and Schaefer [HS13] extended the L_0 minimization framework from images to meshes to gradually remove noise while preserving sharp features. In this paper, we investigate how the L_0 gradient minimization can be applied for smoothing homogenous materials. For the noise on the feature boundaries, we show how to employ this method to retain and sharpen boundaries from blurred boundaries after a local filter for smoothing boundaries. Furthermore, the proposed method is much faster than the topology-controlled method of Günther et al. [GJR*14]; just several seconds for a 128^3 volume.

3. Volume Denoising / Smoothing Using L_0 Gradient Minimization

During the volume acquisition procedure, noise and artifacts are introduced by a number of factors, such as capturing instruments, data quantization and transmission, and discrete sources of radiation [MGM*04]. For example, in medical imaging, additive Gaussian noise usually corrupts data from CT [GBG04] and low intensity MRI [BL03], multiplicative Poisson noise often appears in data from functional MRI [HZPS01], and multiplicative Speckle noise commonly degrades data from Ultrasound [Bur78]. Boundaries of objects usually assume a sharp, discontinuous change in the measured physical value. However, capturing instruments are typically band-limited with a Gaussian frequency response. Thus, the measured boundaries are usually blurred by the Gaussian filter and this causes artifacts in boundaries [KD98].

In this paper, we assume that the scalar volume can be roughly separated into two parts: regions of relatively homogeneous material and boundaries between different materials. For noise in regions of relatively homogeneous material, we introduce the L_0 volume gradient minimization method to suppress it due to its relatively low gradient. This globally retains salient boundaries and regions with nearly constant values. On the other hand, this method cannot effectively remove the noise on the boundaries, as these noisy boundaries usually have higher gradient magnitudes than voxels in relatively homogeneous regions. To address this problem, we further propose the blurring-sharpening strategy, which firstly applies a local filter to blur noise, especially noise on the boundaries, and then employs the L_0 volume gradient minimization method to further remove noise in both materials and boundaries, and to narrow and retain the sharp boundaries from the blurred boundaries after the local filter.

In this section, we will first briefly review the L_0 gradient minimization for a 1D signal, then extend this framework to smooth 3D volume with global low-amplitude structures removal and salient boundaries preservation, and finally describe the alternating optimization strategy to quickly solve this minimization problem. The blurring-sharpening strategy will be presented in Section 4.

3.1. 1D L_0 Smoothing

Let $x = (x_1, x_2, ..., x_n)$ denote the input 1D discrete signal. The L_0 norm of a vector is the number of non-zero values, and it discretely measures the information sparsity (i.e. the smaller the L_0 norm is, the more sparse is the vector), as follows:

$$c(x) = \sharp\{i||x_i| \neq 0\}, \tag{1}$$

where i is the index, and $\sharp\{.\}$ denotes the operator counting the number of x_i that satisfies $|x_i| \neq 0$.

The L_0 gradient norm of the vector is defined to count the number of non-zero value changes of x_i, as follows:

$$c(\nabla x) = \sharp\{i||(x_{i+1} - x_{i-1})/2| \neq 0\}, \tag{2}$$

where $i+1$ and $i-1$ are indexes of neighboring values. $|(x_{i+1} - x_{i-1})/2|$ is the gradient magnitude of one value at

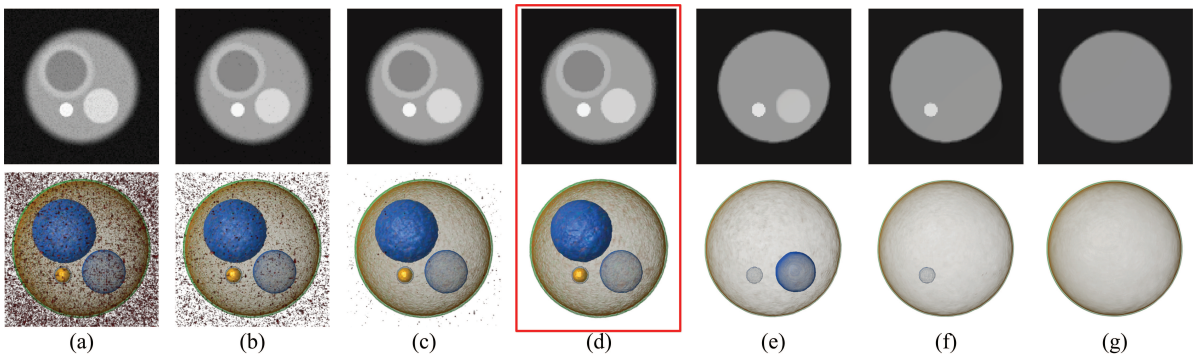

Figure 2: *The impact of the smoothing parameter λ of the L_0 volume gradient minimization method. The slices from the corresponding volumes are displayed at the top row. (a) The original sphere volume with severe artificial noise. (b-g) The smoothed results by the L_0 volume gradient minimization method with the increasing values of λ. Our method obtains the best smoothing result in (d) when $\lambda = 0.002$, since a small λ dose not filter noise sufficiently and large λ destroys the salient spheres.*

the ith position, approximated as the central difference between the corresponding values. Equation 2 is the L_0 norm of the gradient magnitude and it only depends on whether the signal alters its contrast or not, regardless of the magnitude of changes.

For 1D smoothing, we first require that the result y should be structurally similar to the original x. Meanwhile, the details of the 1D signal, which usually have small gradient magnitudes, should be smoothed, i.e. the result y should have a small L_0 norm value. Thus, the objective minimization function for 1D L_0 smoothing is defined as follows:

$$\min_y \{(y-x)^2 + \lambda c(\nabla y)\}, \tag{3}$$

where λ is a smoothing parameter to control the significance of $c(\nabla y)$. The first part of Equation 3 tries to preserve the overall shape as much as possible, which means the overall difference between two vectors should be minimized. The second part of Equation 3 controls the sparsity of the smoothed result. Thus, to minimize the total cost, the significant edges should be retained, otherwise the cost will raise due to the first part of Equation 3. Different from previous edge-aware methods that may destroy the salient edges in different scales due to the local filtering, the L_0 smoothing avoids blurriness of the salient edges and even sharpens these salient edges because of the discrete counting and preservation of the large gradient magnitude values. The smoothing parameter λ has a great influence on the smoothing effect. A large λ results in few significant edges, but it still retains the most sharp edges, which characterizes the dominant information about the 1D signal.

Figure 1(c) shows an example to illustrate the 1D L_0 smoothing effect. It is a 1D profile plot of the scalar values taken from a slice as indicated by the white line in the 2D images from Figure 1(a) and (b). It is clear that the smoothed curve suppresses the details and retains the high-contrast edges.

3.2. 3D Volume L_0 Smoothing

Based on 1D L_0 smoothing above, we can extend it to 3D volume smoothing. Let u denote the input 3D volume and g the unknown smoothed result. The gradient $\nabla g_i = (g_{i,x}, g_{i,y}, g_{i,z})^T$ is calculated from the scalar value difference between neighboring voxels along the x, y and z directions for the ith voxel of the smoothed result. Other sophisticated methods [HAM11] can better estimate the gradient. However, this simple measure is sufficient for our method in most cases. Our gradient measure is expressed as follows:

$$c(\nabla g) = \sharp\{i | \sqrt{(g_{i,x})^2 + (g_{i,y})^2 + (g_{i,z})^2} \neq 0\}, \tag{4}$$

where $\sharp\{.\}$ is the same as Equation 1, and it is the L_0 norm of gradients of the smoothed volume. This operator counts how many voxels there are whose gradient magnitude is not zero in the smoothed volume. It looks for the gradient information sparsity, indicating the number of non-zero gradients.

Similar to 1D signal smoothing, 3D volume smoothing also seeks a balance between result similarity and structure flattening with the input volume, and the 3D volume L_0 gradient minimization equation is defined as follows:

$$\min_g \{(g-u)^2 + \lambda c(\nabla g)\}, \tag{5}$$

where the parameter λ is used to effectively control the level of structure coarseness, and can be adjusted by the user. A large λ results in few boundaries while a small one preserves the boundaries as much as possible.

As shown in Figure 1, the original box volume with two concentric boxes is noise-free. We add synthetic noise (Gaussian noise with $\mu = 0$ and $\sigma = 0.01$) in Figure 1(a).

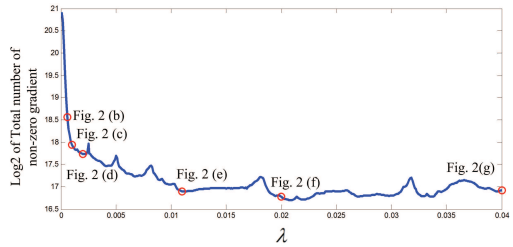

Figure 3: *Number of voxels with a non-zero gradient on the sphere volume in Figure 2 under different λ values. With the increasing of the parameter, more voxels with a non-zero gradient are smoothed. The red circles on the curve indicate the corresponding parameters in Figure 2(b-f), respectively.*

Our denoised result is presented in Figure 1(b), flattening the noise while containing the most salient feature boundaries.

Figure 2 shows the effect of different values of λ on a sphere volume data. The original volume consists of several spheres, whose feature boundaries and inner regions are contaminated by severe noise. A small λ cannot filter out the noise in Figure 2(b-c) and therefore cannot convey the benefits of smoothing. Increasing the value of λ gradually removes noise in regions of homogeneous material while retaining important features. More structures, such as the spheres on the top, are gradually smoothed and only the largest sphere with the largest gradient is preserved (Figure 2(d-g)). Previous local-filtering based methods usually smooth small features first. However, our method smoothes less significant features first, i.e. features with less contrast with their surrounding materials. Thus, the second largest sphere on the top left corner disappears first because it is closer to the surrounding material in terms of intensity and its gradient is smaller than other spheres.

Figure 3 shows the number of voxels with a non-zero gradient on the sphere volume when increasing the λ value from 0 to 0.04, and it indicates the relation between the smoothing effect and the λ value. The first local minimum position (the red circle denoted as Fig. 2(d) on the curve) shows that small gradient details or noise have already been filtered out. With the increase λ, spheres similar with their background will be gradually blurred and finally smoothed. During this process, the number of voxels with a non-zero gradient will first increase and then decrease, due to the balance of the two energy terms in Equation 5. In this example, according to this curve, we choose the optimal value $\lambda = 0.002$ to obtain the best smoothing result in Figure 2(d).

3.3. Computation using Alternative Minimization

The L_0 volume gradient minimization is difficult to optimize due to its non-convex and non-derivative nature, and traditional optimization techniques such as the gradient descent are no longer applicable. Recently, Xu et al. [XLXJ11] and He and Schaefer [HS13] introduced a novel numerical method to minimize the L_0 norm. It is an alternating optimization strategy with the half-quadratic splitting scheme, by introducing a set of auxiliary variables. We briefly describe how to solve the 3D L_0 volume gradient minimization problem.

A set of auxiliary variables δ is firstly introduced in the 3D L_0 volume gradient minimization problem, and Equation 5 becomes:

$$\min_{g,\delta}\{(g-u)^2+\beta(\nabla g-\delta)^2+\lambda c(\delta)\}. \quad (6)$$

The problem above can be solved by alternative minimization:

$$\delta^{(k+1)} \leftarrow \underset{\delta}{argmin}\{\beta(\nabla g^{(k)}-\delta^{(k)})^2+\lambda c(\delta^{(k)})\}, \quad (7)$$

$$g^{(k+1)} \leftarrow \underset{g}{argmin}\{(g^{(k)}-u)^2+\beta(\nabla g^{(k)}-\delta^{(k+1)})^2\}, \quad (8)$$

where k is the iteration number. The parameter β is starting from a small value, and multiplied by 2 each iteration. Equation 6 approaches Equation 5 when β is large enough. Calculating $\delta^{(k+1)}$ corresponds to a shrinkage problem while calculating $g^{(k+1)}$ corresponds to the screened Poisson equation [BCCZ08]. To minimize Equation 7, if $\sqrt{\frac{\lambda}{\beta}} > \nabla g_i^{(k)}$, $\delta_i^{(k+1)} = 0$, otherwise $\delta_i^{(k+1)} = \nabla g_i^{(k)}$, where $g_i^{(k)}$ is the ith voxel in the kth iteration. The Euler-Lagrange solution of Equation 8 is defined:

$$g-\beta\nabla^2 g = u-\beta\nabla\cdot\delta^{(k+1)}, \quad (9)$$

where ∇^2 is the Laplace-Beltrami operator and $\nabla\cdot$ is the divergence operator. Equation 9 can be solved either using the 3D Fast Fourier Transform (FFT) or the preconditioned conjugate gradient method with an appropriate preconditioner [Saa03]. In this paper, we use the FFT solver for its efficiency:

$$g^{(k+1)} \leftarrow F^{-1}\left(\frac{F(u-\beta div(\delta^{k+1}))}{1-\beta lap}\right), \quad (10)$$

where F is the FFT, div is the discrete version of divergence operator $\nabla\cdot$ and lap is the optical transfer function of the discrete Laplacian filter that is evaluated only once. This special alternating optimization strategy is time consuming when the number of iterations needed is large. For example, it takes 52 iterations and about 37.8s to produce Figure 2(d). In our implementation, the accelerated iterative shrinkage algorithm [BYA13] based on first order proximal operators [PB14] introduces an efficient warm-start solution to improve the computational efficiency. Please refer to the paper [BYA13] for the accelerated iterative shrinkage algorithm. To produce the similar result of Figure 2(d), the computation time is 3.95s with this acceleration, nearly 9 times faster than the original alternating optimization method.

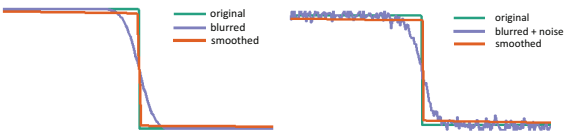

Figure 4: *Boundary smoothing using the blurring-sharpening strategy. (a) To mimic the effect of the band-limited property of capturing instruments, a Gaussian filter is applied to the original 1D sharp signal and this obtains the purple curve with a blurred edge. The red curve is the smoothed result of the L_0 gradient minimization. (b) The captured signal (the purple curve) is further destroyed by noise. The red curve is the recovered signal by the blurring-sharpening strategy.*

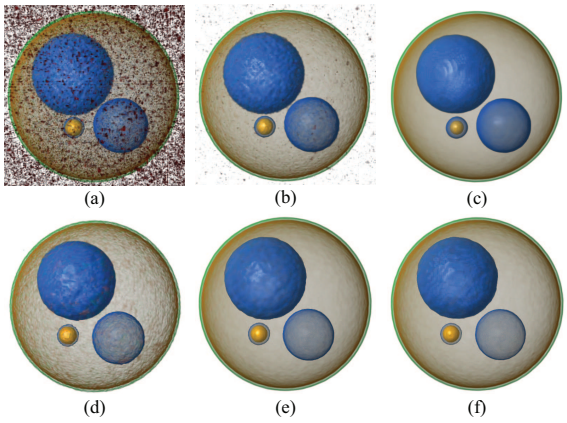

Figure 5: *The impact of the Gaussian filtering in the blurring-sharpening strategy. (a) The original sphere volume with severe noise. (b) The under-smoothed result of the Gaussian filter with $\sigma = 0.5$. (c) The over-smoothed result of the Gaussian filter with $\sigma = 2.0$. (d-f) are smoothed results of (a-c) using the L_0 volume gradient minimization.*

4. Boundary Smoothing Using Blurring-Sharpening Strategy

The L_0 volume gradient minimization method suppresses noise in homogeneous materials and preserves local details on the boundaries between different materials. If boundaries also contain noise, this noise is more visible in the rendered results. However, the proposed method cannot filter out the noise due to its large gradient magnitude, and this results in artifacts or discontinuities in the smoothed result.

In order to remove noise or smooth local details on the boundaries, we propose a blurring-sharpening strategy, which can remove noise both in homogeneous materials and on the boundaries. This strategy first applies a filter to smooth noise in homogeneous materials and on the boundaries. Similar to the Gaussian frequency response of capturing instruments, we use the Gaussian filter in our implementation.

The local filter also blurs the boundaries. We utilize the L_0 volume gradient minimization method in Section 3 to achieve volume smoothing and boundary sharpening. The minimization equation becomes

$$\min_g \{(g - u_{blur})^2 + \lambda c(\nabla g)\}, \tag{11}$$

where u_{blur} is the blurred volume generated by the local filter. The L_0 volume gradient minimization method further filters out noise in homogeneous materials and on the boundaries, and more importantly, it sharpens the blurred boundaries and recovers the original boundaries as much as possible. Thus, noise in homogeneous materials and on boundaries is removed, while boundaries are well preserved or even enhanced.

Figure 4 demonstrates the boundary recovering effect using the blurring-sharpening strategy. The green line in Figure 4 represents a sharp, discontinuous edge. The purple line is obtained by blurring the edge with the Gaussian filter in Figure 4(a), and the line is further degraded by noise in Figure 4(b). The blurred purple line is applying the L_0 gradient minimization method, and the edge is sharpened and recovered to the red line similar to the original green line in Figure 4(a). The blurring-sharpening strategy is used to recover the blurred and noisy purple line in Figure 4(b), and it obtains the red line similar to the original green line. Thus, the proposed strategy not only smoothes the signal, but also narrows the blurred signal, recovers the sharpness, and retains the signal as far as possible.

As shown in Figure 5, noise on the boundaries results in visual artifacts in volume visualization, such as the small bumps in the L_0 smoothed volume in Figure 5(d). This is because boundaries have larger gradient magnitudes than those in other regions. Therefore, these boundaries together with the noise are preserved after the L_0 volume gradient minimization method. In order to suppress noise on the boundaries, the blurring-sharpening strategy is applied. After the Gaussian filter with different smoothing capabilities, the under-smoothed result in Figure 5(b) still contains noise, and the over-smoothed result in Figure 5(c) blurs the important features, i.e. the yellow sphere becomes smaller and the distance between two blue spheres gets shorter. The blurring-sharpening strategy can both faithfully recover the original boundaries in Figure 5(e) and (f). These results show that our blurring-sharpening strategy suppresses noise not only in homogeneous materials but also on the boundaries, and is more robust with respect to the Gaussian parameter.

5. Results and Discussion

We have implemented the proposed volume smoothing method with C++ and Matlab. Several volume datasets in d-

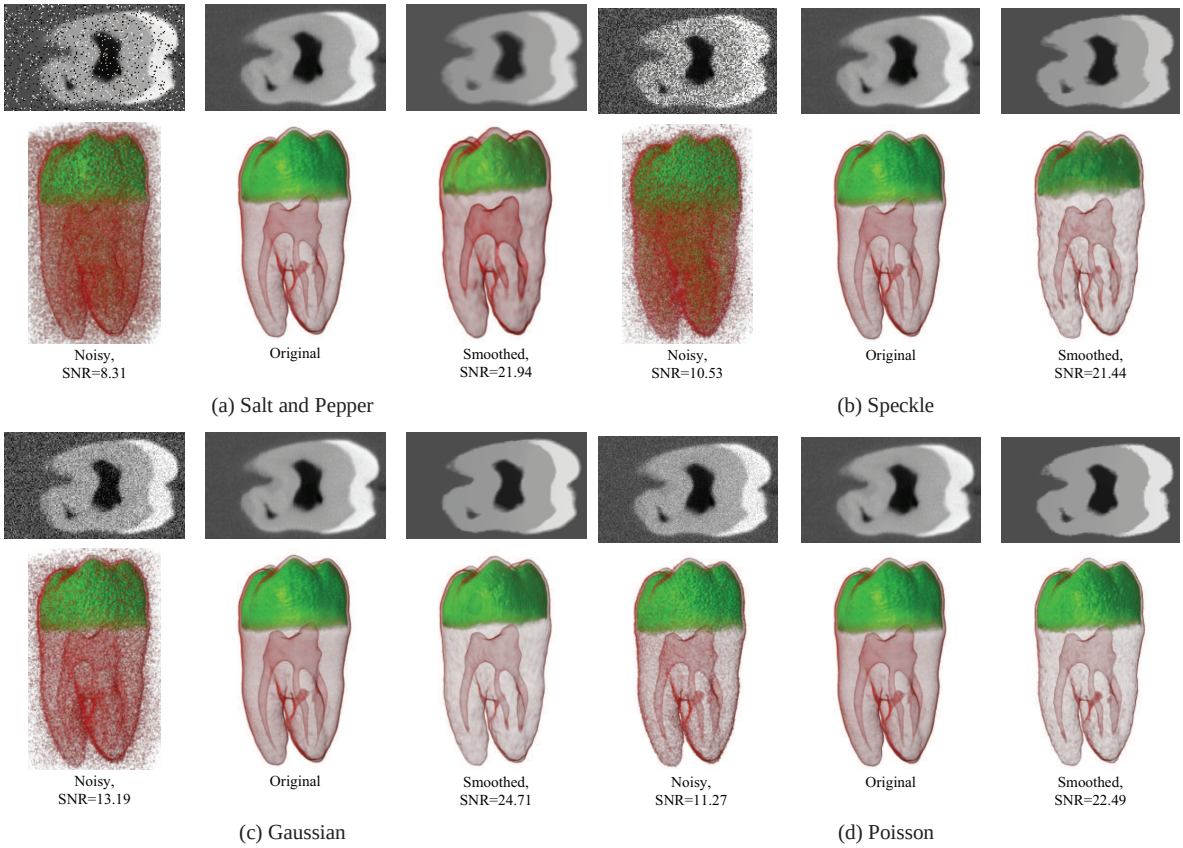

Figure 6: *The denoising capability of the blurring-sharpening strategy based on the tooth volume. (a-d) are obtained by adding one particular type of noise, as indicated by the corresponding captions. SNR (in dB) of the noisy and the smoothed volumes are shown in each figure.*

ifferent modalities were used to verify the effectiveness and efficiency of our method in volume denoising / smoothing. We performed all experiments on a PC with a dual core 3.0GHz and 8GB memory.

Four common noises, Salt and Pepper noise, multiplicative Speckle noise, additive Gaussian noise, and additive Poisson noise, are added to a nearly noise-free tooth data (CT) to test the denoising capability of our method. Figure 6 shows the denoising effect of our method in all four cases. These results demonstrate that our method can effectively remove common noises generated from capturing instruments while preserving significant features. For example, additive Gaussian noise introduces gaussian-distributed type artifacts in the volume, and the blurring-sharpening strategy not only removes noise in homogeneous material effectively, such as enamel, dentin and pulp, but also suppresses noise on the boundaries, such as cementum.

A 3D MRI Tomato data is used to compare three volume smoothing methods: the bilateral filter, edge aware anisotropic diffusion [HM10], and the proposed L_0 volume gradient minimization method. This data has a high signal-to-noise (SNR = 17.09), and the noise distributed over the whole volume makes boundaries of interesting features difficult to recognize, as shown in Figure 7(a). The bilateral filter cannot remove noise as indicated in the blue box of Figure 7(b), and blurs small features, and even removes the small seeds in the red box due to its local filtering and averaging operation. Edge aware anisotropic diffusion preserves peripheral boundaries and diminishes noise out of the tomato. However, it blurs boundaries of internal features, for example the noisy seeds and placenta in the red box of Figure 7(c), as noisy structures of these features make the evaluation of the second derivative not very accurate. Our method globally suppresses local details with small gradient magnitudes, producing visual smoothness in homogeneous materials, and at the same time retains and sharpens prominent boundaries, such as the boundaries of seeds and placenta as indicated in Figure 7(d). Our method is more effective in smoothing noisy regions in this volume.

We further compare our method with the bilateral filter

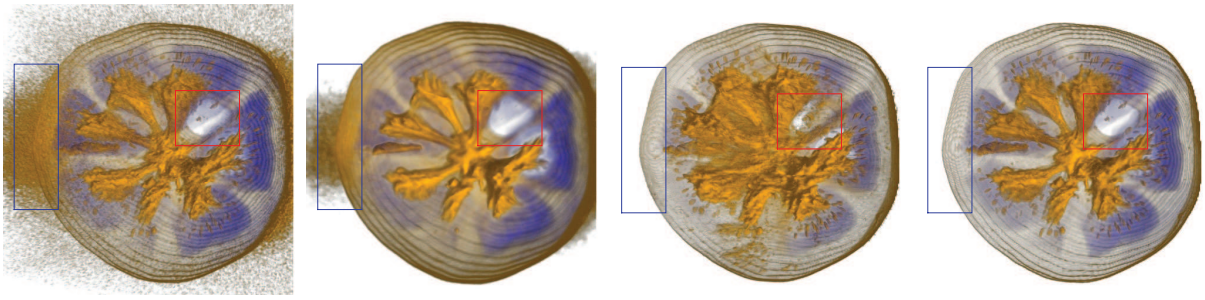

(a) (b) (c) (d)

Figure 7: *Comparison of three smoothing methods on a 3D MRI tomato volume. (a) The original volume data. (b) The smoothed volume with the bilateral filter, $\sigma_s = 2$, $\sigma_r = 40$. (c) The smoothed volume with edge aware anisotropic diffusion, $\sigma = 2$ and 10 iterations. (d) Our smoothed result, $\lambda = 0.00023$. Our method suppresses low gradient amplitude details and globally retains salient edges, clearly shown in the boundaries of seeds and placenta.*

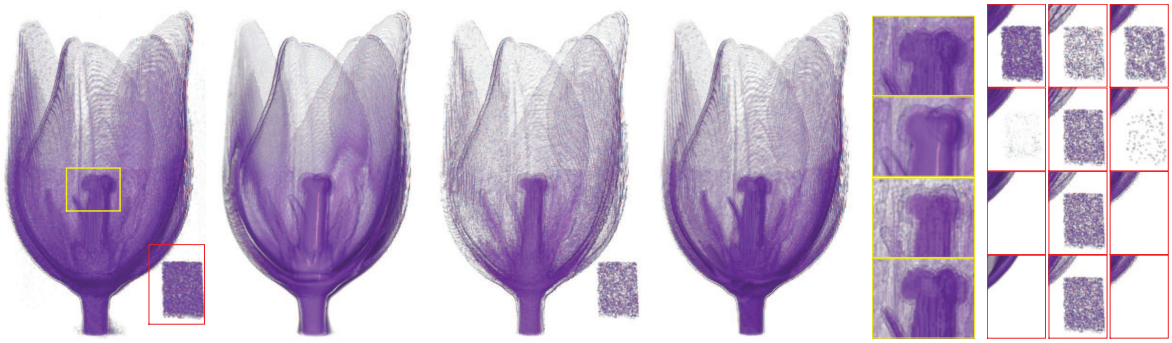

(a) (b) (c) (d) (e) (f)

Figure 8: *Comparison of three smoothing methods on a 3D CT flower volume. (a) The original volume data. (b) The smoothed volume with the bilateral filter, $\sigma_s = 5$, $\sigma_r = 120$. (c) The smoothed volume with edge aware anisotropic diffusion, $\sigma = 10$ and 15 iterations. (d) Our L_0 smoothing result, $\lambda = 0.005$. Compared with other methods, our method suppresses low gradient amplitude block-like noises, while retaining the salient boundaries. (e) Close-ups (as shown the yellow box in (a)) of pistils and stamens in the results. (f) Close-ups (as shown the red box in (a)) of the results with increasing smoothing strength gradually from top to bottom. The methods from left to right are the bilateral filter, edge aware anisotropic diffusion and the L_0 volume gradient minimization method.*

and edge aware anisotropic diffusion [HM10] by smoothing area noise in a flower volume [IYYI14]. This volume was artificially contaminated with a block noise near the stem, as shown in the red box of Figure 8(a). Figure 8(e) shows the pistils and stamens in the results with three smoothing methods. Figure 8(f) shows the results when each method increases the smoothing capability. The bilateral filter can remove the block noise, but also blurs the boundaries of pistils and stamens due to its averaging operation, as shown in Figure 8(b). Edge aware anisotropic diffusion cannot quite smooth out the block noise in Figure 8(c) due to the relatively large second directive in the block noise. When increasing the λ value, our method gradually filters out the block noise while still preserving the flower's boundary as shown in Figure 8(f). Our method produces a smooth and continuous boundary while suppressing the low-contrast structures, as shown in Figure 8(d). Compared with the bilateral filter and edge aware anisotropic diffusion, our method can effectively deal with this type of area noise.

In Figure 9, we present our smoothed result of 3D ultrasound data of human liver [RPT]. Ultrasound volumes are usually ruined with Speckle noise, and features are often low-contrast. It is extremely difficult to visually detect important structures, as shown in Figure 9(a). Our method filters out Speckle noise around important features in this volume, and retains and sharpens salient boundaries. Tubular structures, which are hardly detectable in Figure 9(a), are more visually distinct and recognizable in Figure 9(b).

Table 1 reports the total processing time (in seconds) of

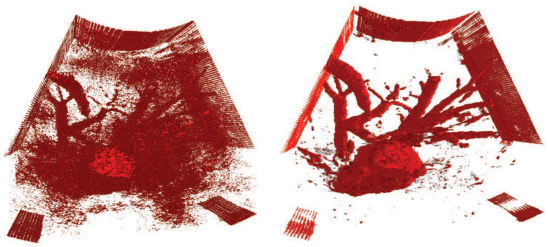

(a) (b)

Figure 9: *(a) The original 3D ultrasound data of human liver. (b) Our smoothed result. Tubular structures are more visually recognizable.*

Table 1: *The computational times (in seconds) of bilateral filter (BLF), edge aware anisotropic diffusion (AD), and our method for Tomato volume (Figure 7), Flower volume (Figure 8), and Human Liver volume (Figure 9).*

Data sets	Resolution	Time (seconds)		
		BLF	AD	Ours
Tomato	$256 \times 256 \times 164$	148.70	229.13	26.78
Flower	$160 \times 168 \times 243$	186.12	145.41	33.21
Human Liver	$493 \times 415 \times 172$	435.26	1388.54	78.55

three methods on different volume datasets. The FFT (Equation 10) is computed in Matlab using MEX files (C/C++ extension for Matlab), and other codes, including the bilateral filter and edge aware anisotropic diffusion, are implemented in C++. As can be seen from Table 1, our method performs about 5 and 10 times faster than the bilateral filter and edge aware anisotropic diffusion, respectively. The main reason is that FFT makes the time of one iteration much faster due to its efficiency, and the accelerated iterative shrinkage greatly reduces the iteration number compared with edge aware anisotropic diffusion. The efficiency of the bilateral filter greatly depends on the kernel size. In our experiment, the parameters σ_s and σ_r are chosen in the range of $[2,5]$ and $[40,120]$, respectively. The iteration number in edge aware anisotropic diffusion also affects the time, which is between 10 and 20 in our experiment.

Table 2 summarizes the performance of three methods in terms of Mean Squared Error (MSE), Structural Similarity Index (SSIM) [WBSS04] and Quality Index Based on Local Variance (QILV) [AFSAW06] for quality assessment. For the experimented tomato and flower volume, our method performs better than both other methods, because it removes the local details with a low gradient magnitude, and globally preserves and even enhances the salient boundaries.

Table 2: *Performance evaluation of different smoothing methods: Bilateral filter (BLF), edge aware anisotropic diffusion (AD), and our method on Tomato volume (Figure 7) and Flower volume (Figure 8). The lower value of MSE is better, while the higher value of SSIM and QILV is better.*

	Tomato			Flower		
	MSE	SSIM	QILV	MSE	SSIM	QILV
BLF	2.20	0.93	0.99	4.69	0.91	0.99
AD	4.24	0.87	0.92	8.58	0.90	0.97
Ours	2.02	0.95	0.99	3.86	0.93	0.99

6. Conclusions

In this paper, we have presented a volume denoising / smoothing method based on the L_0 gradient minimization framework for volume visualization. The basic idea is to globally control the number of voxels of a non-zero gradient in the result. It suppresses less important details and noise with a low gradient magnitude, and sharpens and enhances salient features. In addition to smooth noise in homogenous materials, we also proposed the blurring-sharpening strategy to diminish noise on the boundaries. This strategy can be used for volume denoising and smoothing effectively, and the result contains salient features with smooth boundaries and visually pleasing structures while being in line with the original volume. Our experiments demonstrated the effectiveness and efficiency of our method. Compared with the bilateral filter and edge aware anisotropic diffusion, our method is more effective in dealing with area noise with a low gradient amplitude.

Future work. Although the L_0 norm does not rely on the magnitude, the L_0 volume gradient minimization method does depend on the gradient magnitude, as it needs to balance the result similarity and structure sparsity. Thus, low-contrast features will be first smoothed by our method, while high-contrast noise cannot be effectively filtered out due to its large gradient magnitude. For future work, we plan to extend our method by designing different types of the gradient magnitude for removing noise adaptively. For example, it can filter out high-contrast noise first. Currently, users heuristically select the optimal value of λ. We aim to find an automatic scheme to suggest the optimal parameter for the input volume. In addition, the alternative minimization strategy is only an approximate solution, which may result in small position differences of surfaces. We would like to find a more efficient and accurate way to solve the minimization problem, such as the fused coordinate descent framework [CZL14]. As the current optimization method could not perform interactive volume smoothing, we would like to also improve the optimization method to denoise / smooth large volumes interactively.

Acknowledgement

The authors would like to thank the anonymous reviewers for their valuable comments. This work was partially supported by 863 Program Project 2012AA12A404 and National Natural Science Foundation of China No. 61472354.

References

[AFSAW06] AJA-FERNANDEZ S., SAN JOSE ESTEPAR R., ALBEROLA LOPEZ C., WESTIN C.-F.: Image quality assessment based on local variance. In *28th IEEE EMBS* (New York City, NY, USA, September 2006), pp. 4815–4818. 9

[BCCZ08] BHAT P., CURLESS B., COHEN M., ZITNICK C. L.: Fourier analysis of the 2d screened poisson equation for gradient domain problems. In *Proceedings of the 10th European Conference on Computer Vision: Part II* (2008), ECCV '08, pp. 114–128. 5

[BCM05] BUADES A., COLL B., MOREL J. M.: A review of image denoising algorithms, with a new one. *Simul 4* (2005), 490–530. 2

[BL03] BAO P., LEI Z.: Noise reduction for magnetic resonance images via adaptive multiscale products thresholding. *IEEE Transactions on Medical Imaging 22*, 9 (2003), 1089–1099. 3

[Bur78] BURCKHARDT C.: Speckle in ultrasound b-mode scans. *IEEE Transactions on Sonics and Ultrasonics 25*, 1 (Jan 1978), 1–6. 3

[BYA13] BADRI H., YAHIA H., ABOUTAJDINE D.: Fast multiscale detail decomposition via accelerated iterative shrinkage. In *SIGGRAPH Asia 2013 Technical Briefs* (2013), SA '13, pp. 33:1–33:4. 5

[CZL14] CHENG X., ZENG M., LIU X.: Feature-preserving filtering with l_0 gradient minimization. *Computers & Graphics 38* (2014), 150–157. 9

[FDCO03] FLEISHMAN S., DRORI I., COHEN-OR D.: Bilateral mesh denoising. *ACM Transactions on Graphics 22*, 3 (July 2003), 950–953. 2

[FFLS08] FARBMAN Z., FATTAL R., LISCHINSKI D., SZELISKI R.: Edge-preserving decompositions for multi-scale tone and detail manipulation. In *ACM SIGGRAPH 2008 Papers* (2008), SIGGRAPH '08, pp. 67:1–67:10. 2

[GBG04] GRADVEL P., BEAUDOIN G., GUISE J. A. D.: A method for modeling noise in medical images. *IEEE Transactions on Medical Imaging 23*, 10 (2004), 1221–1232. 3

[GJR*14] GÜNTHER D., JACOBSON A., REININGHAUS J., SEIDEL H.-P., SORKINE-HORNUNG O., WEINKAUF T.: Fast and memory-efficient topological denoising of 2D and 3D scalar fields. *IEEE Transactions on Visualization and Computer Graphics (Proceedings Scientific Visualization 2014) 20*, 12 (December 2014). 2, 3

[HAM11] HOSSAIN Z., ALIM U. R., MÖLLER T.: Toward high-quality gradient estimation on regular lattices. *IEEE Transactions on Visualization and Computer Graphics 17*, 4 (Apr. 2011), 426–439. 4

[HM10] HOSSAIN Z., MÖLLER T.: Edge aware anisotropic diffusion for 3D scalar data. *IEEE Transactions on Visualization and Computer Graphics (Proceedings Scientific Visualization 2010) 16*, 6 (Nov-Dec 2010), 1375–1384. 1, 2, 7, 8

[HS13] HE L., SCHAEFER S.: Mesh denoising via l0 minimization. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH 2013) 32*, 4 (July 2013), 64:1–64:8. 3, 5

[HZPS01] HAGBERG G. E., ZITO G., PATRIA F., SANES J. N.: Improved detection of event-related functional mri signals using probability functions. *NeuroImage 2* (2001), 1193–1205. 3

[IYYI14] IJIRI T., YOSHIZAWA S., YOKOTA H., IGARASHI T.: Flower modeling via x-ray computed tomography. *ACM Transactions on Graphics 33*, 4 (July 2014), 48:1–48:10. 8

[JWS12] JACOBSON A., WEINKAUF T., SORKINE O.: Smooth shape-aware functions with controlled extrema. *Computer Graphics Forum (Proceedings of EUROGRAPHICS/ACM SIGGRAPH Symposium on Geometry Processing) 31*, 5 (2012), 1577–1586. 2

[KAF09] KRISSIAN K., AJA-FERNÁNDEZ S.: Noise-driven anisotropic diffusion filtering of mri. *IEEE Transactions on Image Processing 18*, 10 (Oct. 2009), 2265–2274. 2

[KD98] KINDLMANN G., DURKIN J. W.: Semi-automatic generation of transfer functions for direct volume rendering. In *Proceedings of the 1998 IEEE Symposium on Volume Visualization* (1998), VVS '98, pp. 79–86. 2, 3

[KMA97] KRISSIAN K., MALANDAIN G., AYACHE N.: Directional anisotropic diffusion applied to segmentation of vessels in 3d images. In *Proceedings of the First International Conference on Scale-Space Theory in Computer Vision* (1997), SCALE-SPACE '97, pp. 345–348. 2

[MAD14] M. AUBRY S. PARIS S. H. J. K., DURAND F.: Fast local laplacian filters: Theory and applications. *ACM Transactions on Graphics* (2014). 2

[MGM*04] MOTWANI M. C., GADIYA M. C., MOTWANI R. C., HARRIS F. C., JR.: Survey of image denoising techniques, 2004. 3

[MMK*98] MÖLLER T., MUELLER K., KURZION Y., MACHIRAJU R., YAGEL R.: Design of accurate and smooth filters for function and derivative reconstruction. In *Proceedings of the 1998 IEEE Symposium on Volume Visualization* (1998), VVS '98, pp. 143–151. 2

[MMMY97] MÖLLER T., MACHIRAJU R., MUELLER K., YAGEL R.: Evaluation and design of filters using a taylor series expansion. *IEEE Transactions on Visualization and Computer Graphics 3*, 2 (Apr. 1997), 184–199. 2

[PB14] PARIKH N., BOYD S.: Proximal algorithms. *Foundations and Trends in Optimization 1*, 3 (2014). 5

[PM90] PERONA P., MALIK J.: Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence 12*, 7 (July 1990), 629–639. 1, 2

[RPT] RICHARD PRAGER A. G., TREECE G.: 3d ultrasound using stradx 7.4g. 8

[Saa03] SAAD Y.: *Iterative Methods for Sparse Linear Systems*, 2nd ed. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2003. 5

[SH05] SIGG C., HADWIGER M.: Fast third-order texture filtering. In *GPU Gems 2*, Pharr M., (Ed.). Addison-Wesley, 2005, pp. 313–329. 2

[TM98] TOMASI C., MANDUCHI R.: Bilateral filtering for gray and color images. In *Proceedings of the Sixth International Conference on Computer Vision* (1998), ICCV '98, pp. 839–. 1, 2

[WBSS04] WANG Z., BOVIK A. C., SHEIKH H. R., SIMONCELLI E. P.: Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing 13*, 4 (2004), 600–612. 9

[WGS10] WEINKAUF T., GINGOLD Y., SORKINE O.: Topology-based smoothing of 2D scalar fields with C1-continuity. *Computer Graphics Forum (Proceedings of EuroVis) 29*, 3 (2010), 1221–1230. 2

[XLXJ11] XU L., LU C., XU Y., JIA J.: Image smoothing via l0 gradient minimization. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH Asia 2011)* (2011). 1, 3, 5