

A Generalized Framework for Edge-Preserving and Structure-Preserving Image Smoothing

Wei Liu,^{1,2} Pingping Zhang,³ Yinjie Lei,^{4*} Xiaolin Huang,^{1,5} Jie Yang,^{1,5*} Ian Reid²

¹Department of Automation, Shanghai Jiao Tong University, ²The University of Adelaide

³Dalian University of Technology, ⁴Sichuan University, ⁵Institute of Medical Robotics, Shanghai Jiao Tong University
{wei.liu02, ian.reid}@adelaide.edu.au, jssxzhpp@mail.dlut.edu.cn, yinjie@scu.edu.cn, {xiaolinhuang, jieyang}@sjtu.edu.cn

Abstract

Image smoothing is a fundamental procedure in applications of both computer vision and graphics. The required smoothing properties can be different or even contradictory among different tasks. Nevertheless, the inherent smoothing nature of one smoothing operator is usually fixed and thus cannot meet the various requirements of different applications. In this paper, a non-convex non-smooth optimization framework is proposed to achieve diverse smoothing natures where even contradictory smoothing behaviors can be achieved. To this end, we first introduce the truncated Huber penalty function which has seldom been used in image smoothing. A robust framework is then proposed. When combined with the strong flexibility of the truncated Huber penalty function, our framework is capable of a range of applications and can outperform the state-of-the-art approaches in several tasks. In addition, an efficient numerical solution is provided and its convergence is theoretically guaranteed even the optimization framework is non-convex and non-smooth. The effectiveness and superior performance of our approach are validated through comprehensive experimental results in a range of applications.

Introduction

The key challenge of many tasks in both computer vision and graphics can be attributed to image smoothing. At the same time, the required smoothing properties can vary dramatically for different tasks. In this paper, depending on the required smoothing properties, we roughly classify a large number of applications into four groups.

Applications in the first group require the smoothing operator to smooth out small details while preserving strong edges, and the amplitudes of these strong edges can be reduced but the edges should be neither blurred nor sharpened. Representatives in this group are image detail enhancement and HDR tone mapping (Farbman et al. 2008; Fattal, Agrawala, and Rusinkiewicz 2007; He, Sun, and Tang 2013). Blurring edges can result in halos while sharpening edges will lead to gradient reversals (Farbman et al. 2008).

*Jie Yang and Yinjie Lei are the corresponding authors of this paper.

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Our method is capable of (a) image detail enhancement, (b) clip-art compression artifacts removal, (c) guided depth map upsampling and (d) image texture removal. These applications are representatives of edge-preserving and structure-preserving image smoothing and require contradictory smoothing properties.

The second group includes tasks like clip-art compression artifacts removal (Nguyen and Brown 2015; Xu et al. 2011), image abstraction and pencil sketch production (Xu et al. 2011). In contrast to the ones in the first group, these tasks require to smooth out small details while sharpening strong edges. This is because edges can be blurred in the compressed clip-art image and they need to be sharpened when the image is recovered (see Fig. 1(b) for example). Sharper edges can produce better visual quality in image abstraction and pencil sketch. At the same time, the amplitudes of strong edges are not allowed to be reduced in these tasks.

Guided image filtering, such as guided depth map upsampling (Park et al. 2011; Ferstl et al. 2013; Liu et al. 2017b) and flash/no flash filtering (Kopf et al. 2007; Petschnigg et al. 2004), is categorized into the third group. The structure inconsistency between the guidance image and target image, which can cause blurring edges and texture copy artifacts in the smoothed image (Ham, Cho, and Ponce 2015; Liu et al. 2017b), should be properly handled by the specially designed smoothing operator. They also need to sharpen edges in the smoothed image due to the reason that low-quality capture of depth and noise in the no flash images can lead to blurred edge (see Fig. 1(c) for example).

Tasks in the fourth group require to smooth the image in a scale-aware manner, e.g., image texture removal (Xu et al. 2012; Zhang et al. 2014; Cho et al. 2014). This kind of tasks require to smooth out small structures even when they contain strong edges, while large structure should be properly

preserved even the edges are weak (see Fig. 1(d) for example). This is totally different from that in the above three groups where they all aim at preserving strong edges.

To be more explicit, we categorize the smoothing procedures in the first to the third groups as *edge-preserving image smoothing* since they try to preserve salient edges, while the smoothing processes in the fourth group are classified as *structure-preserving image smoothing* because they aim at preserving salient structures.

A diversity of edge-preserving and structure-preserving smoothing operators have been proposed for various tasks. Generally, each of them is designed to meet the requirements of certain applications, and thus its inherent smoothing nature is usually fixed. Therefore, there is seldom a smoothing operator that can meet all the smoothing requirements of the above four groups, which are quite different or even contradictory. For example, the L_0 norm smoothing (Xu et al. 2011) can sharpen strong edges and is suitable for clip-art compression artifacts removal, however, this will lead to gradient reversals in image detail enhancement and HDR tone mapping. The weighted least squares (WLS) smoothing (Farbman et al. 2008) performs well in image detail enhancement and HDR tone mapping, but it is not capable of sharpening edges and structure-preserving smoothing, etc.

In contrast to most of the smoothing operators in the literature, a new smoothing operator, which is based on a non-convex non-smooth optimization framework, is proposed in this paper. It can achieve different and even contradictory smoothing behaviors and is able to handle the applications in the four groups mentioned above. The main contributions of this paper are as follows:

1. We introduce the *truncated Huber penalty* function which has seldom been used in image smoothing. By varying the parameters, it shows strong flexibility.
2. A robust non-convex non-smooth optimization framework is proposed. When combined with the strong flexibility of the truncated Huber penalty function, our model can achieve various and even contradictory smoothing behaviors. We show that it is able to handle the tasks in the four groups mentioned above. This has seldom been achieved by previous smoothing operators.
3. An efficient numerical solution to the proposed optimization framework is provided. Its convergence is theoretically guaranteed.
4. Our method is able to outperform the specially designed approaches in many tasks and state-of-the-art performance is achieved.

Related Work

Tremendous smoothing operators have been proposed in recent decades. In terms of edge-preserving smoothing, bilateral filter (BLF) (Tomasi and Manduchi 1998) is the early work that has been used in various tasks such as image detail enhancement (Fattal, Agrawala, and Rusinkiewicz 2007), HDR tone mapping (Durand and Dorsey 2002), etc. However, it is prone to produce results with gradient reversals and halos (Farbman et al. 2008). Its alternatives (Gastal and

Oliveira 2012; 2011) also share a similar problem. Guided filter (GF) (He, Sun, and Tang 2013) can produce results free of gradient reversals but halos still exist. The WLS smoothing (Farbman et al. 2008) solves a global optimization problem and performs well in handling these artifacts. The L_0 norm smoothing is able to eliminate low-amplitude structures while sharpening strong edges, which can be applied to the tasks in the second group. To handle the structure inconsistency problem, Shen et al. (Shen et al. 2015b) proposed to perform mutual-structure joint filtering. They also explored the relation between the guidance image and target image via optimizing a scale map (Shen et al. 2015a), however, additional processing was adopted for structure inconsistency handling. Ham et al. (Ham, Cho, and Ponce 2015) proposed to handle the structure inconsistency by combining a static guidance weight with a Welsch's penalty (Holland and Welsch 1977) regularized smoothness term, which led to a static/dynamic (SD) filter. Gu et al. (Gu et al. 2017b) presented a weighted analysis representation model for guided depth map enhancement.

In terms of structure-preserving smoothing, Zhang et al. (Zhang et al. 2014) proposed to smooth structures of different scales with a rolling guidance filter (RGF). Cho et al. (Cho et al. 2014) modified the original BLF with local patch-based analysis of texture features and obtained a bilateral texture filter (BTF) for image texture removal. Karacan et al. (Karacan, Erdem, and Erdem 2013) proposed to smooth image textures by making use of region covariances that captured local structure and textural information. Xu et al. (Xu et al. 2012) adopted the relative total variation (RTV) as a prior to regularize the texture smoothing procedure. Fan et al. (Fan et al. 2018; 2019) proposed to perform various kinds of image smoothing through convolutional neural networks. Chen et al. (Chan and Esedoglu 2005) proved that the TV- L_1 model (Chan and Esedoglu 2005; Nikolova 2004) could smooth images in a scale-aware manner, and it is thus ideal for structure-preserving smoothing such as image texture removal (Buades et al. 2010).

Most of the approaches mentioned above are limited to a few applications because their inherent smoothing natures are usually fixed. In contrast, our method proposed in this paper can have strong flexibility in achieving various smoothing behaviors, which enables wider applications of our method than most of them. Moreover, our method can show better performance than these methods in several applications that they are specially designed for.

Our Approach

Truncated Huber Penalty Function

We first introduce the truncated Huber penalty function which is defined as:

$$h_T(x) = \begin{cases} h(x), & |x| \leq b \\ b - \frac{a}{2}, & |x| > b \end{cases} \quad \text{s.t. } a \leq b, \quad (1)$$

where a, b are constants. $h(\cdot)$ is the Huber penalty function (Huber and others 1964) defined as:

$$h(x) = \begin{cases} \frac{1}{2a}x^2, & |x| < a \\ |x| - \frac{a}{2}, & |x| \geq a \end{cases}, \quad (2)$$




Figure 2: Plots of (a) different penalty functions and (b) the truncated Huber penalty function with different parameter settings.

$h_T(\cdot)$ and $h(\cdot)$ are plotted in Fig. 2(a) with $a = \epsilon$ which is a sufficient small value (e.g., $\epsilon = 10^{-7}$). $h(\cdot)$ is an edge-preserving penalty function, but it cannot sharpen edges when adopted to regularize the smoothing procedure. In contrast, $h_T(\cdot)$ can sharpen edges because it is able to not penalize image edges due to the truncation. The Welsch's penalty function (Holland and Welsch 1977), which was adopted in the recent proposed SD filter (Ham, Cho, and Ponce 2015), is also plotted in the figure. This penalty function is known to be capable of sharpening edges, which is also because it seldom penalizes strong image edges. The Welsch's penalty function is close to the L_2 norm when the input is small, while the $h_T(\cdot)$ can be close to the L_1 norm when a is set sufficient small, which demonstrates $h_T(\cdot)$ can better preserve weak edges than the Welsch's penalty function.

With different parameter settings, $h_T(\cdot)$ can show strong flexibility to yield different penalty behaviors. Assume the input intensity values are within $[0, I_m]$, then the amplitude of any edge will fall in $[0, I_m]$. We first set $a = \epsilon$. Then if we set $b > I_m$, $h_T(\cdot)$ will be actually the same as $h(\cdot)$ because the second condition in Eq. (1) can never be met. Because a is sufficient small, $h_T(\cdot)$ will be close to the L_1 norm in this case, and thus it will be an edge-preserving penalty function that does not sharpen edges. Conversely, when we set $b < I_m$, the truncation in $h_T(\cdot)$ will be activated. This can lead to having penalization on weak edges without penalizing strong edges, and thus the strong edges are sharpened. To be short, b can act as a switch to decide whether $h_T(\cdot)$ can sharpen edges or not. Similarly, by setting $a = b > I_m$ and $a = b < I_m$, $h_T(\cdot)$ can be easily switched between the L_2 norm and truncated L_2 norm. Note that the truncated L_2 norm is also able to sharpen edges (Xu, Zheng, and Jia 2013). In contrast, the Welsch's penalty function does not enjoy this kind of flexibility. Different cases of $h_T(\cdot)$ are illustrated in Fig. 2(b).

Model

Given an input image f and a guidance image g , the smoothed output image u is the solution to the following objective function:

$$E_u(u) = \sum_i \left(\sum_{j \in N_d(i)} h_T(u_i - f_j) + \lambda \sum_{j \in N_s(i)} \omega_{i,j} h_T(u_i - u_j) \right), \quad (3)$$

where h_T is defined in Eq.(1); $N_d(i)$ is the $(2r_d + 1) \times (2r_d + 1)$ square patch centered at i ; $N_s(i)$ is the $(2r_s + 1) \times (2r_s + 1)$ square patch centered at i ; λ is a parameter that controls the overall smoothing strength. To be clear, we adopt $\{a_d, b_d\}$ and $\{a_s, b_s\}$ to denote the parameters of $h_T(\cdot)$ in the data term and smoothness term, respectively. The guidance weight $\omega_{i,j}$ is defined as:

$$\omega_{i,j} = \frac{1}{(|g_i - g_j| + \delta)^\alpha}, \quad (4)$$

where α determines the sensitivity to the edges in g which can be the input image, i.e., $g = f$. $|\cdot|$ represents the absolute value. δ is a small constant being set as $\delta = 10^{-7}$.

The adoption of $h_T(\cdot)$ makes our model in Eq. (3) to enjoy a strong flexibility. As will be shown in the following property analysis section, with different parameter settings, our model is able to achieve different smoothing behaviors, and thus it is capable of various tasks that require either edge-preserving smoothing or structure-preserving smoothing.

Numerical Solution

Our model in Eq. (3) is not only non-convex but also non-smooth, which arises from the adopted $h_T(\cdot)$. Commonly used approaches (Lanczkiet and Sriperumbudur 2009; Nikolova and Ng 2005; Wang et al. 2008; Zhang, Kwok, and Yeung 2004) for solving non-convex optimization problems are not applicable. To tackle this problem, we first rewrite $h_T(\cdot)$ in a new equivalent form. By defining $\nabla_{i,j}^d = u_i - f_j$ and $\nabla_{i,j}^s = u_i - u_j$, we have:

$$h_T(\nabla_{i,j}^*) = \min_{l_{i,j}^*} \left\{ h(\nabla_{i,j}^* - l_{i,j}^*) + (b_* - \frac{a_*}{2}) |l_{i,j}^*|_0 \right\}, \quad (5)$$

where $*$ $\in \{d, s\}$, $|l_{i,j}^*|_0$ is the L_0 norm of $l_{i,j}^*$. The minimum of the right side of Eq. (5) is obtained on the condition:

$$l_{i,j}^* = \begin{cases} 0, & |\nabla_{i,j}^*| \leq b_* \\ \nabla_{i,j}^*, & |\nabla_{i,j}^*| > b_* \end{cases}, \quad *, * \in \{d, s\}. \quad (6)$$

The detailed proof of Eq. (5) and Eq. (6) is provided in our supplementary file. These two equations also theoretically validate our analysis in Fig. 2(b): we have $|\nabla_{i,j}^*| \in [0, I_m]$ if the intensity values are in $[0, I_m]$. Then if $b > I_m$, based on Eq. (5) and Eq. (6), we will always have $h_T(\nabla_{i,j}^*) = h(\nabla_{i,j}^*)$ which means $h_T(\cdot)$ degrades to $h(\cdot)$.

A new energy function is defined as:

$$E_{ul}(u, l^d, l^s) = \sum_{i,j} \left(h(\nabla_{i,j}^d - l_{i,j}^d) + (b_d - \frac{a_d}{2}) |l_{i,j}^d|_0 \right) + \lambda \sum_{i,j} \omega_{i,j} \left(h(\nabla_{i,j}^s - l_{i,j}^s) + (b_s - \frac{a_s}{2}) |l_{i,j}^s|_0 \right). \quad (7)$$

Based on Eq. (5) and Eq. (6), we then have:

$$E_u(u) = \min_{l^*} E_{ul}(u, l^d, l^s), \quad *, * \in \{d, s\}. \quad (8)$$

Given Eq. (6) as the optimum condition of Eq. (8) with respect to l^* , optimizing $E_{ul}(u, l^d, l^s)$ with respect to u only involves Huber penalty function $h(\cdot)$. The problem can thus be optimized through the half-quadratic (HQ) optimization technique (Geman and Yang 1995; Nikolova and Ng 2005). More specifically, a variable μ^* ($*$ $\in \{d, s\}$) and a function $\psi(\mu_{i,j}^*)$ with respect to μ^* exist such that:

$$h(\nabla_{i,j}^* - l_{i,j}^*) = \min_{\mu_{i,j}^*} \left\{ \mu_{i,j}^* (\nabla_{i,j}^* - l_{i,j}^*)^2 + \psi(\mu_{i,j}^*) \right\}, \quad (9)$$

where the optimum is yielded on the condition:

$$\mu_{i,j}^* = \begin{cases} \frac{1}{2a_*}, & |\nabla_{i,j}^* - l_{i,j}^*| < a_* \\ \frac{1}{2|\nabla_{i,j}^* - l_{i,j}^*|}, & |\nabla_{i,j}^* - l_{i,j}^*| \geq a_* \end{cases}, \quad * \in \{d, s\}. \quad (10)$$

The detailed proof of Eq. (9) and Eq. (10) is provided in our supplementary file. Then we can further define a new energy function:

$$\begin{aligned} E_{ul\mu}(u, l^d, l^s, \mu^d, \mu^s) = & \sum_{i,j} \left(\mu_{i,j}^d (\nabla_{i,j}^d - l_{i,j}^d)^2 + \psi(\mu_{i,j}^d) + (b_d - \frac{a_d}{2}) |l_{i,j}^d|_0 \right) + \\ & \lambda \sum_{i,j} \omega_{i,j} \left(\mu_{i,j}^s (\nabla_{i,j}^s - l_{i,j}^s)^2 + \psi(\mu_{i,j}^s) + (b_s - \frac{a_s}{2}) |l_{i,j}^s|_0 \right). \end{aligned} \quad (11)$$

Based on Eq. (9) and Eq. (10), we then have:

$$E_{ul}(u, l^*) = \min_{\mu^*} E_{ul\mu}(u, l^*, \mu^*), \quad * \in \{d, s\}. \quad (12)$$

Given Eq. (10) as the optimum condition of μ^* for Eq. (12), optimizing $E_{ul\mu}(u, l^d, l^s, \mu^d, \mu^s)$ with respect to u only involves the L_2 norm penalty function, which has a closed-form solution. However, since the optimum conditions in Eq. (6) and Eq. (10) both involve u , therefore, the final solution u can only be obtained in an iterative manner. Assuming we have got u^k , then $(l^*)^k$ and $(\mu^*)^k$, $(* \in \{s, d\})$ can be updated through Eq. (6) and Eq. (10) with u^k . Finally, u^{k+1} is obtained with:

$$u^{k+1} = \min_u E_{ul\mu}(u, (l^*)^k, (\mu^*)^k), \quad (13)$$

Eq.(13) has a close-form solution as:

$$u^{k+1} = \left(A^k - 2\lambda \mathcal{W}^k \right)^{-1} \left(D^k + 2\lambda S^k \right), \quad (14)$$

where \mathcal{W}^k is an affinity matrix with $\mathcal{W}_{i,j}^k = \omega_{i,j} (\mu_{i,j}^s)^k$, A^k is a diagonal matrix with $A_{ii}^k = \sum_{j \in N_d(i)} (\mu_{i,j}^d)^k + 2\lambda \sum_{j \in N_s(i)} \omega_{i,j} (\mu_{i,j}^s)^k$, D^k is a vector with $D_i^k = \sum_{j \in N_d(i)} (\mu_{i,j}^d)^k (f_j + (l_{i,j}^d)^k)$ and S^k is also a vector with $S_i^k = \sum_{j \in N_s(i)} \omega_{i,j} (\mu_{i,j}^s)^k (l_{i,j}^s)^k$.

The above optimization procedure monotonically decreases the value of $E_u(u)$ in each step, its convergence is theoretically guaranteed. Given u^k in the k th iteration and $* \in \{s, d\}$, then for any u , we have:

$$E_u(u) \leq E_{ul}(u, (l^*)^k), \quad E_u(u^k) = E_{ul}(u^k, (l^*)^k), \quad (15)$$

$$\begin{cases} E_{ul}(u, (l^*)^k) \leq E_{ul\mu}(u, (l^*)^k, (\mu^*)^k) \\ E_{ul}(u^k, (l^*)^k) = E_{ul\mu}(u^k, (l^*)^k, (\mu^*)^k) \end{cases}. \quad (16)$$

Given $(l^*)^k$ has been updated through Eq. (6), Eq. (15) is based on Eq. (8) and Eq. (5). After $(\mu^*)^k$ has been updated through Eq. (10), Eq. (16) is based on Eq. (12) and Eq. (9). We now have:

$$\begin{aligned} E_{ul}(u^{k+1}, (l^*)^k) & \leq E_{ul\mu}(u^{k+1}, (l^*)^k, (\mu^*)^k) \\ & \leq E_{ul\mu}(u^k, (l^*)^k, (\mu^*)^k) = E_{ul}(u^k, (l^*)^k), \end{aligned} \quad (17)$$

the first and second inequalities follow from Eq. (16) and Eq. (13), respectively. We finally have:

$$E_u(u^{k+1}) \leq E_{ul}(u^{k+1}, (l^*)^k) \leq E_{ul}(u^k, (l^*)^k) = E_u(u^k), \quad (18)$$

Algorithm 1 Image Smoothing via Non-convex Non-smooth Optimization

Require: Input image f , guide image g , iteration number N , parameter $\lambda, \alpha, a_*, b_*, r_*, u^0 \leftarrow f$, with $* \in \{d, s\}$

1: **for** $k = 0 : N$ **do**

2: With u^k , compute $(\nabla_{i,j}^*)^k$, update $(l_{i,j}^*)^k$ according to Eq. (6)

3: With $(l_{i,j}^*)^k$, update $(\mu_{i,j}^*)^k$ according to Eq. (10)

4: With $(l_{i,j}^*)^k$ and $(\mu_{i,j}^*)^k$, solve for u^{k+1} according to Eq. (13) (or Eq. (14))

5: **end for**

Ensure: Smoothed image u^{N+1}

the first and second inequalities follow from Eq. (15) and Eq. (17), respectively. Since the value of $E_u(u)$ is bounded from below, Eq. (18) indicates that the convergence of our iterative scheme is theoretically guaranteed.

The above optimization procedure is iteratively performed N times to get the output u^N . In all our experiments, we set $u^0 = f$, which is able to produce promising results in each application. Our optimization procedure is summarized in Algorithm 1.

Property Analysis

With different parameter settings, the strong flexibility of $h_T(\cdot)$ makes our model able to achieve various smoothing behaviors. First, we show that some classical approaches can be viewed as special cases of our model. For example, by setting $a_d = b_d > I_m, a_s = \epsilon, b_s > I_m, \alpha = 0, r_d = 0, r_s = 1$, our model is an approximation of the TV model (Rudin, Osher, and Fatemi 1992) which is a representative edge-preserving smoothing operator. If we set $\alpha = 0.2, g = f$ with other parameters the same as above, then the first iteration of Algorithm 1 will be the WLS smoothing (Farbman et al. 2008) which performs well in handling gradient reversals and halos in image detail enhancement and HDR tone mapping. With parameters $a_d = \epsilon, b_d > I_m, a_s = \epsilon, b_s > I_m, \alpha = 0, r_d = 0, r_s = 1$, our model can yield very close smoothing natures as the TV- L_1 model (Buades et al. 2010) which is classical for structure-preserving smoothing.

For different kinds of applications, our model can produce better results than the special cases mentioned above. To be convenient, we first start with the tasks in the fourth group which require structure-preserving smoothing. For these tasks, the parameters are set as $a_d = \epsilon, b_d > I_m, a_s = \epsilon, b_s > I_m, r_d = r_s, \alpha = 0.5, g = f$. This parameter setting has the following two advantages: first, the setting $a_d = \epsilon, b_d > I_m, a_s = \epsilon, b_s > I_m$ enables our model to have the structure-preserving property similar to that of the TV- L_1 model; second, the guidance weight with $\alpha = 0.5, g = f$ can make our model to obtain sharper edges in the results than the TV- L_1 model does. We illustrate this with 1D smoothing results in Fig. 3(a) and (b). Fig. 6(b) and (c) further show a comparison of image texture removal results. As shown in the figure, both the TV- L_1 model and our model can properly remove the small textures, however, edges in our result are much sharper than that in the result of




Figure 3: 1D signal with structures of different scales and amplitudes. Smoothing result of (a) TV- L_1 smoothing (Buades et al. 2010), (c) WLS (Farbman et al. 2008), (e) SD filter (Ham, Cho, and Ponce 2015), our results in (b), (d) and (f).




Figure 4: Image detail enhancement results of different approaches. (a) Input image. Result of (b) WLS (Farbman et al. 2008) and (c) our method. The upper parts of each close-up in (b) and (c) correspond to the patches in the smoothed image.




Figure 5: Clip-art compression artifacts removal results of different approaches. (a) Input image. (b) Our result. Close-ups of (c) input image and results of (d) SD filter (Ham, Cho, and Ponce 2015), (e) our method with the structure-preserving parameter setting, (f) our method with the edge-preserving and structure-preserving parameter setting.




Figure 6: Texture smoothing results of different approaches. (a) Input image. Result of (b) TV- L_1 smoothing (Buades et al. 2010), and (c) our method.

the TV- L_1 model. The typical values for $r_d = r_s$ are $1 \sim 3$ depending on the texture size. λ is usually smaller than 1. Larger r_d, r_s, λ can lead larger structures to be removed. The iteration number is set as $N = 10$.

When dealing with image detail enhancement and HDR tone mapping in the first group, one way is to set the parameters so that our model can perform WLS smoothing. In contrast, we can further make use of the structure-preserving property of our model to produce better results. The parameters are set as follows: $a_d = \epsilon, b_d > I_m, a_s = \epsilon, b_s > I_m, r_d = r_s, \alpha = 0.2, g = f$. This kind of parameter setting is based on the following observation in our experiments: when we adopt $N = 1$ and set λ to a large value, the amplitudes of different structures will decrease at different rates, i.e., the amplitudes of small structures can have a

larger decrease than the large ones, as illustrated in Fig. 3(d). At the same time, edges are neither blurred nor sharpened. This kind of smoothing behavior is desirable for image detail enhancement and HDR tone mapping. As a comparison, Fig. 3(c) shows the smoothing result of the WLS smoothing. As can be observed from the figures, our method can better preserve the edges (see the bottom of the 1D signals in Fig. 3(c) and (d)). Fig. 4(b) and (c) further show a comparison of image detail enhancement results. We fix $r_d = r_s = 2$ and vary λ to control the smoothing strength. λ for the tasks in the first group is usually much larger than that for the ones in the fourth group, for example, the result in Fig. 4(c) is generated with $\lambda = 20$.

To sharpen edges that is required by the tasks in the second and the third groups, we can set $b_s < I_m$ in the smoothness term. In addition, we further set other parameters as $a_d = \epsilon, b_d < I_m, a_s = \epsilon$. The truncation $b_d < I_m$ in the data term can help our model to be robust against the outliers in the input image, for example, the noise in the no flash image and low-quality depth map. The truncation $b_s < I_m$ in the smoothness term can enable our model to be an edge-preserving one. By setting $a_d = a_s = \epsilon$, our model can further enjoy the structure-preserving property. With both edge-preserving and structure-preserving smoothing natures, our model has the ability to preserve large structures with weak edges and small structures with strong edges at the same




Figure 7: HDR tone mapping results of different approaches. Result of (a) BF (Tomasi and Manduchi 1998), (b) GF (He, Sun, and Tang 2013), (c) L_0 norm smoothing (Xu et al. 2011), (d) WLS (Farbman et al. 2008), (e) SG-WLS (Liu et al. 2017a) and (f) our method.




Figure 8: Clip-art compression artifacts removal results of different methods. (a) Input compressed image. Result of (b) the approach proposed by Wang et al. (Wang, Wong, and Heng 2006), (c) L_0 norm smoothing (Xu et al. 2011), (d) region fusion approach (Nguyen and Brown 2015), (e) BTF (Cho et al. 2014) and (f) our method.

time, which is challenging but is of practical importance. Fig. 5(a) illustrates this kind of case with an example of clip-art compression artifacts removal: both the thin black circle around the “wheel” and the gray part in the center of the “wheel” should be preserved. The challenge lies on two facts. On one hand, if we perform edge-preserving smoothing, the gray part will be removed because the corresponding edge is weak. Fig. 5(d) shows the result of the SD filter (Ham, Cho, and Ponce 2015). The SD filter can properly preserve the thin black circle and sharpen the edges thanks to the adopted Welsch’s penalty function, however, it fails to preserve the weak edge between the black part and the gray part. On the other hand, if we adopt structure-preserving smoothing, then the thin black circle will be smoothed due to its small structure size. Fig. 5(e) shows the corresponding result of our method with the structure-preserving parameter setting described above. In contrast, our method with the edge-preserving and structure-preserving parameter setting can preserve both these two parts and sharpen the edges, as shown in Fig. 5(f). Fig. 3(e) and (f) also show a comparison of the SD filter and our method with 1D smoothing results. We fix $\alpha = 0.5$, $r_d = r_s$, $N = 10$ for the tasks in both the second and the third groups. We empirically set $b_d = b_s = 0.05I_m \sim 0.2I_m$ and $r_d = r_s = 1 \sim 5$ depending on the applied task and the input noise level.

The structure inconsistency issue in the third group can also be easily handled by our model. Note that $\mu_{i,j}^s$ in Eq. (11) is computed with the smoothed image in each iteration, as formulated in Eq. (10), it thus can reflect the inherent natures of the smoothed image. The guidance weight

$\omega_{i,j}$ can provide additional structural information from the guidance image g . This means that $\mu_{i,j}^s$ and $\omega_{i,j}$ can complement each other. In fact, the equivalent guidance weight of Eq. (11) in each iteration is $\mu_{i,j}^s \omega_{i,j}$, which can reflect the property of both the smoothed image and the guidance image. In this way, it can properly handle the structure inconsistency problem and avoid blurring edges and texture copy artifacts. Similar ideas were also adopted in (Ham, Cho, and Ponce 2015; Liu et al. 2017b).

Applications and Experimental Results

Our method is applied to various tasks in the first to the fourth groups to validate the effectiveness. Comparisons with the state-of-the-art approaches in each application are also presented. Due to the limited space, we only show experimental results of four applications.

Our experiments are performed on a PC with an Intel Core i5 3.4GHz CPU (one thread used) and 8GB memory. For an RGB image of size 800×600 and $N = 10$ in Algorithm 1, the running time is 10.04/25.09/43.11/69.82/96.73 seconds in MATLAB for $r_d = r_s = 1/2/3/4/5$. Note that as described in the property analysis section, the value of $r_d = r_s$ is smaller than 3 in most cases except for guided depth map upsampling. For the tasks in the first group which require $N = 1$, the computational cost could be further reduced to $\frac{1}{10}$ of that mentioned above.

HDR tone mapping is a representative task in the first group. It requires to decompose the input image into a base layer and a detail layer through edge-preserving smoothing.

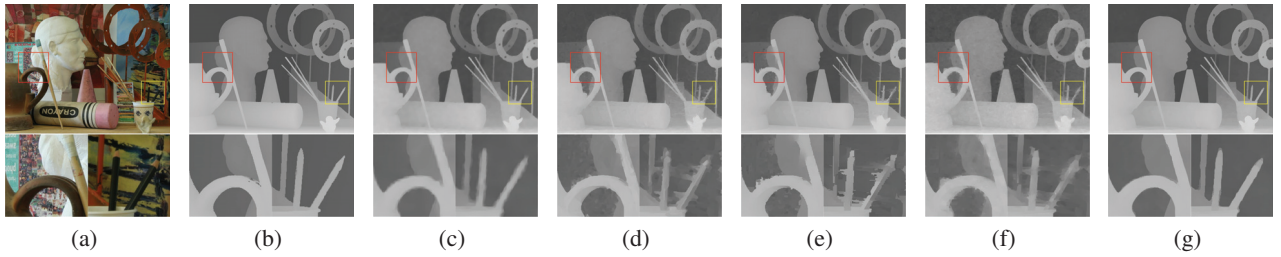


Figure 9: Guided depth map upsampling results of simulated ToF data. (a) Guidance color image. (b) Ground-truth depth map. Result of (c) the approach proposed by Gu et al. (Gu et al. 2017b), (d) SGF (Zhang et al. 2015), (e) SD filter (Ham, Cho, and Ponce 2015), (f) Park et al. (Park et al. 2011) and (g) our method.




Figure 10: Guided depth upsampling results of real ToF data. (a) Guidance intensity image. (b) Ground-truth depth map. Result of (c) the approach proposed by Gu et al. (Gu et al. 2017b), (d) TGV (Ferstl et al. 2013), (e) SD filter (Ham, Cho, and Ponce 2015), (f) SGF (Zhang et al. 2015) and (g) our method.




Figure 11: Image texture removal results. (a) Input image. Result of (b) JCAS (Gu et al. 2017a), (c) RTV (Xu et al. 2012), (d) FCN based approach (Chen, Xu, and Koltun 2017), (e) muGIF (Guo et al. 2018) (f) BTF (Cho et al. 2014) and (g) our method.

The challenge of this task is that if the edges are sharpened by the smoothing procedure, it will result in gradient reversals, and halos will occur if the edges are blurred. Fig. 7 shows the tone mapping results using different edge-preserving smoothing operators. The results of BF (Tomasi and Manduchi 1998) and GF (He, Sun, and Tang 2013) contain clear halos around the picture frames and the light fixture, as shown in Fig. 7(a) and (b). This is due to their local smoothing natures where strong smoothing can also blur salient edges (Farbman et al. 2008; He, Sun, and Tang 2013). The L_0 norm smoothing (Xu et al. 2011) can properly eliminate halos, but there are gradient reversals in its result as illustrated in Fig. 7(c). This is because the L_0 smoothing is prone to sharpen salient edges. The WLS (Farbman et al. 2008) and SG-WLS (Liu et al. 2017a) smoothing perform well in handling gradient reversals and halos in most cases. However, there are slight halos in their results as illustrated

in the left close-up in Fig. 7(d) and (e). These artifacts are properly eliminated in our results.

Clip-art compression artifacts removal. Clip-art images are piecewise constant with sharp edges. When they are compressed in JPEG format with low quality, there will be edge-related artifacts, and the edges are usually blurred as shown in Fig. 8(a). Therefore, when removing the compression artifacts, the edges should also be sharpened in the restored image. We thus classify this task into the second group. The approach proposed by Wang et al. (Wang, Wong, and Heng 2006) can seldom handle heavy compression artifacts as shown in Fig. 8(b). The L_0 norm smoothing fails to preserve weak edges as shown in Fig. 8(c). The region fusion approach (Nguyen and Brown 2015) is able to produce results with sharpened edges, however, it also enhances the blocky artifacts along strong edges as highlighted in Fig. 8(d). The edges in the result of BTF (Cho et al. 2014)

Table 1: Quantitative comparison on the noisy simulated ToF data. Results are evaluated in MAE. The best results are in **bold**. The second best results are underlined.

	<i>Art</i>				<i>Book</i>				<i>Dolls</i>				<i>Laundry</i>				<i>Moebius</i>				<i>Reindeer</i>			
	2×	4×	8×	16×	2×	4×	8×	16×	2×	4×	8×	16×	2×	4×	8×	16×	2×	4×	8×	16×	2×	4×	8×	16×
TGV(Ferstl et al. 2013)	0.8	<u>1.21</u>	2.01	4.59	0.61	0.88	1.21	2.19	0.66	0.95	<u>1.38</u>	2.88	<u>0.61</u>	0.87	<u>1.36</u>	3.06	0.57	<u>0.77</u>	1.23	2.74	0.61	0.85	1.3	3.41
AR(Yang et al. 2014)	1.17	1.7	2.93	5.32	0.98	1.22	1.74	2.89	0.97	1.21	1.71	2.74	1	1.31	1.97	3.43	0.95	1.2	1.79	2.82	1.07	1.3	2.03	3.34
SG-WLS(Liu et al. 2017a)	1.26	1.9	3.07	-	0.82	1.12	1.73	-	0.87	1.11	1.81	-	0.86	1.17	2	-	0.82	1.08	1.79	-	0.9	1.32	2.01	-
FGI(Li et al. 2016b)	0.9	1.37	2.46	4.89	0.66	<u>0.85</u>	<u>1.23</u>	<u>1.96</u>	0.74	<u>0.95</u>	1.41	<u>2.13</u>	0.71	0.99	1.59	2.67	0.67	0.82	<u>1.2</u>	<u>1.87</u>	0.75	0.94	1.55	2.73
SD Filter(Ham, Cho, and Ponce 2015)	1.42	1.85	3.06	5.55	0.84	1.11	1.76	3.03	0.87	1.2	1.88	3.26	0.74	1.1	1.96	3.63	0.81	1.13	1.84	3.16	0.93	1.25	2.03	3.67
FBS(Barron and Poole 2016)	1.16	1.64	2.88	5.52	0.86	1.1	1.57	2.68	1.04	1.27	1.73	2.76	0.96	1.25	1.94	3.54	0.93	1.14	1.68	2.75	1.05	1.31	1.99	3.43
muGIF(Guo et al. 2018)	1.93	2.39	3.29	5.05	1.42	1.55	1.76	2.48	1.33	1.45	1.69	2.26	1.32	1.49	1.77	2.67	1.16	1.29	1.61	2.44	1.63	1.76	2.01	2.69
Park et al.(Park et al. 2011)	1.00	1.26	<u>2.00</u>	<u>3.46</u>	0.73	0.89	1.35	2.15	0.85	1.04	1.50	2.45	0.64	0.87	<u>1.36</u>	<u>2.57</u>	0.67	0.85	1.35	2.25	0.78	0.94	1.39	<u>2.52</u>
Shen et al.(Shen et al. 2015b)	1.66	2.47	3.44	5.55	1.19	1.47	2.06	3.1	1.19	1.56	2.15	3.04	1.34	1.73	2.41	3.85	1.2	1.5	2.13	2.95	1.26	1.65	2.46	3.66
Gu et al.(Gu et al. 2017b)	1.79	2.21	3.2	5.04	1.34	1.69	2.25	3.13	1.37	1.58	2.05	2.85	1.49	1.74	2.34	3.5	1.34	1.56	2.09	2.99	1.29	1.55	2.19	3.33
Li et al.(Li et al. 2016a)	0.61	1.46	2.98	5.09	0.52	0.95	1.87	2.98	<u>0.63</u>	1.02	1.89	2.92	0.58	1.14	2.21	3.58	<u>0.53</u>	0.96	1.89	2.99	0.52	1.07	2.17	3.59
Ours	-	3.77	4.49	6.29	-	3.21	3.28	3.79	-	3.19	3.28	3.79	-	3.34	3.61	4.45	-	3.23	3.35	3.92	-	3.39	3.65	4.54
	<u>0.69</u>	1.07	1.65	2.96	<u>0.55</u>	0.81	1.22	1.78	0.62	0.9	1.27	1.84	<u>0.61</u>	<u>0.89</u>	1.28	2.12	0.51	0.75	1.12	1.71	<u>0.56</u>	<u>0.87</u>	1.27	2.08

Table 2: Quantitative comparison on real ToF dataset. The errors are calculated as MAE to the measured ground-truth in mm. The best results are in **bold**. The second best results are underlined.

	<i>Books</i>	<i>Devil</i>	<i>Shark</i>
Bicubic	16.23mm	17.78mm	16.66mm
GF(He, Sun, and Tang 2013)	15.55mm	16.1mm	17.1mm
SD Filter(Ham, Cho, and Ponce 2015)	13.47mm	15.99mm	16.18mm
SG-WLS(Liu et al. 2017a)	14.71mm	16.24mm	16.51mm
Shen et al.(Shen et al. 2015b)	15.47mm	16.18mm	17.33mm
Park et al.(Park et al. 2011)	14.31mm	15.36mm	15.88mm
TGV(Ferstl et al. 2013)	<u>12.8mm</u>	<u>14.97mm</u>	<u>15.53mm</u>
AR(Yang et al. 2014)	14.37mm	15.41mm	16.27mm
Gu et al.(Gu et al. 2017b)	13.87mm	15.36mm	15.88mm
SGF(Zhang et al. 2015)	13.57mm	15.74mm	16.21mm
FGI(Li et al. 2016b)	14.21mm	16.43mm	16.37mm
FBS(Barron and Poole 2016)	15.93mm	17.21mm	16.33mm
Li et al.(Li et al. 2016a)	14.33mm	15.09mm	15.82mm
Ours	12.49mm	14.51mm	15.02mm

are blurred in Fig. 8(e). Our result is illustrated in Fig. 8(f) with edges sharpened and compression artifacts removed.

Guided depth map upsampling belongs to the guided image filtering in the third group. The RGB guided image can provide additional structural information to restore and sharpen the depth edges. The challenge of this task is the structure inconsistency between the depth map and the RGB guidance image, which can cause blurring depth edges and texture copy artifacts in the upsampled depth map. We test our method on the simulated dataset provided in (Yang et al. 2014). Fig. 9 shows the visual comparison between our result and the results of the recent state-of-the-art approaches. Our method shows better performance in preserving sharp depth edges and avoiding texture copy artifacts. Tab. 1 also shows the quantitative evaluation on the results of different methods. Following the measurement used in (Guo et al. 2018; Li et al. 2016b; Liu et al. 2017a; Yang et al. 2014), the evaluation is measured in terms of mean absolute errors (MAE). As Tab. 1 shows, our method can achieve the best or the second best performance among all the compared approaches.

We further validate our method on the real data introduced by Ferstl et al. (Ferstl et al. 2013). The real dataset contains three low-resolution depth maps captured by a ToF depth camera and the corresponding highly accurate ground-truth depth maps captured with structured light. The upsampling factor for the real dataset is $\sim 6.25\times$. The visual comparison

in Fig. 10 and the quantitative comparison in Tab. 2 shows that our method can outperform the compared methods and achieve state-of-the-art performance.

Image texture removal belongs to the tasks in the fourth group. It aims at extracting salient meaningful structures while removing small complex texture patterns. The challenge of this task is that it requires structure-preserving smoothing rather than the edge-preserving in the above tasks. Fig. 11(a) shows a classical example of image texture removal: the small textures with strong edges should be smoothed out while the salient structures with weak edges should be preserved. Fig. 11(b)~(f) show the results of the recent state-of-the-art approaches. The joint convolutional analysis and synthesis sparse (JCAS) model (Gu et al. 2017a) can well remove the textures, but the resulting edges are also blurred. The RTV method (Xu et al. 2012), muGIF (Guo et al. 2018), BTF (Cho et al. 2014) and FCN based approach (Chen, Xu, and Koltun 2017) cannot completely remove the textures, in addition, the weak edges of the salient structures have also been smoothed out in their results. Our method can both preserve the weak edges of the salient structures and remove the small textures.

Conclusion

We propose a non-convex non-smooth optimization framework for edge-preserving and structure-preserving image smoothing. We first introduce the truncated Huber penalty function which shows strong flexibility. Then a robust framework is presented. When combined with the flexibility of the truncated Huber penalty function, our framework is able to achieve different and even contradictory smoothing behaviors using different parameter settings. This is different from most previous approaches of which the inherent smoothing natures are usually fixed. We further propose an efficient numerical solution to our model and prove its convergence theoretically. Comprehensive experimental results in a number of applications demonstrate the effectiveness of our method.

Acknowledgement

This paper is partly supported by NSFC, China (No. U1803261, 61977046), Key Research and Development Program of Sichuan Province (No. 2019YFG0409) and National Key Research and Development Project (No. 2018AAA0100702)

References

- Barron, J. T., and Poole, B. 2016. The fast bilateral solver. In *ECCV*, 617–632. Springer.
- Buades, A.; Le, T. M.; Morel, J.-M.; Vese, L. A.; et al. 2010. Fast cartoon+ texture image filters. *TIP* 19(8):1978–1986.
- Chan, T. F., and Esedoglu, S. 2005. Aspects of total variation regularized l1 function approximation. *SIAM Journal on Applied Mathematics* 65(5):1817–1837.
- Chen, Q.; Xu, J.; and Koltun, V. 2017. Fast image processing with fully-convolutional networks. In *ICCV*, volume 9, 2516–2525.
- Cho, H.; Lee, H.; Kang, H.; and Lee, S. 2014. Bilateral texture filtering. *ToG* 33(4):128.
- Durand, F., and Dorsey, J. 2002. Fast bilateral filtering for the display of high-dynamic-range images. In *ToG*, volume 21, 257–266. ACM.
- Fan, Q.; Yang, J.; Wipf, D.; Chen, B.; and Tong, X. 2018. Image smoothing via unsupervised learning. In *SIGGRAPH Asia 2018 Technical Papers*, 259. ACM.
- Fan, Q.; Chen, D.; Yuan, L.; Hua, G.; Yu, N.; and Chen, B. 2019. A general decoupled learning framework for parameterized image operators. *IEEE transactions on pattern analysis and machine intelligence*.
- Farbman, Z.; Fattal, R.; Lischinski, D.; and Szeliski, R. 2008. Edge-preserving decompositions for multi-scale tone and detail manipulation. In *ToG*, volume 27, 67. ACM.
- Fattal, R.; Agrawala, M.; and Rusinkiewicz, S. 2007. Multiscale shape and detail enhancement from multi-light image collections. In *ToG*, volume 26, 51. ACM.
- Farstl, D.; Reinbacher, C.; Ranftl, R.; R  ther, M.; and Bischof, H. 2013. Image guided depth upsampling using anisotropic total generalized variation. In *ICCV*, 993–1000.
- Gastal, E. S., and Oliveira, M. M. 2011. Domain transform for edge-aware image and video processing. In *ToG*, volume 30, 69. ACM.
- Gastal, E. S., and Oliveira, M. M. 2012. Adaptive manifolds for real-time high-dimensional filtering. *ToG* 31(4):33.
- Geman, D., and Yang, C. 1995. Nonlinear image recovery with half-quadratic regularization. *TIP* 4(7):932–946.
- Gu, S.; Meng, D.; Zuo, W.; and Zhang, L. 2017a. Joint convolutional analysis and synthesis sparse representation for single image layer separation. In *ICCV*, 1717–1725. IEEE.
- Gu, S.; Zuo, W.; Guo, S.; Chen, Y.; Chen, C.; and Zhang, L. 2017b. Learning dynamic guidance for depth image enhancement. In *CVPR*.
- Guo, X.; Li, Y.; Ma, J.; and Ling, H. 2018. Mutually guided image filtering. *TPAMI*.
- Ham, B.; Cho, M.; and Ponce, J. 2015. Robust image filtering using joint static and dynamic guidance. In *CVPR*, 4823–4831.
- He, K.; Sun, J.; and Tang, X. 2013. Guided image filtering. *TPAMI* 35(6):1397–1409.
- Holland, P. W., and Welsch, R. E. 1977. Robust regression using iteratively reweighted least-squares. *Communications in Statistics-theory and Methods* 6(9):813–827.
- Huber, P. J., et al. 1964. Robust estimation of a location parameter. *The annals of mathematical statistics* 35(1):73–101.
- Karacan, L.; Erdem, E.; and Erdem, A. 2013. Structure-preserving image smoothing via region covariances. *ToG* 32(6):176.
- Kopf, J.; Cohen, M. F.; Lischinski, D.; and Uyttendaele, M. 2007. Joint bilateral upsampling. In *ToG*, volume 26, 96. ACM.
- Lanckriet, G. R., and Sriperumbudur, B. K. 2009. On the convergence of the concave-convex procedure. In *NeurIPS*, 1759–1767.
- Li, Y.; Huang, J.-B.; Ahuja, N.; and Yang, M.-H. 2016a. Deep joint image filtering. In *ECCV*, 154–169. Springer.
- Li, Y.; Min, D.; Do, M. N.; and Lu, J. 2016b. Fast guided global interpolation for depth and motion. In *ECCV*, 717–733. Springer.
- Liu, W.; Chen, X.; Shen, C.; Liu, Z.; and Yang, J. 2017a. Semi-global weighted least squares in image filtering. In *ICCV*, 5861–5869.
- Liu, W.; Chen, X.; Yang, J.; and Wu, Q. 2017b. Robust color guided depth map restoration. *TIP* 26(1):315–327.
- Nguyen, R. M., and Brown, M. S. 2015. Fast and effective l0 gradient minimization by region fusion. In *ICCV*, 208–216.
- Nikolova, M., and Ng, M. K. 2005. Analysis of half-quadratic minimization methods for signal and image recovery. *SIAM Journal on Scientific Computing* 27(3):937–966.
- Nikolova, M. 2004. A variational approach to remove outliers and impulse noise. *Journal of Mathematical Imaging and Vision* 20(1-2):99–120.
- Park, J.; Kim, H.; Tai, Y.-W.; Brown, M. S.; and Kweon, I. 2011. High quality depth map upsampling for 3d-tof cameras. In *ICCV*, 1623–1630. IEEE.
- Petschnigg, G.; Szeliski, R.; Agrawala, M.; Cohen, M.; Hoppe, H.; and Toyama, K. 2004. Digital photography with flash and no-flash image pairs. *ToG* 23(3):664–672.
- Rudin, L. I.; Osher, S.; and Fatemi, E. 1992. Nonlinear total variation based noise removal algorithms. *Physica D: nonlinear phenomena* 60(1-4):259–268.
- Shen, X.; Yan, Q.; Xu, L.; Jia, J.; et al. 2015a. Multispectral joint image restoration via optimizing a scale map. *TPAMI* (1):1–1.
- Shen, X.; Zhou, C.; Xu, L.; and Jia, J. 2015b. Mutual-structure for joint filtering. In *ICCV*, 3406–3414.
- Tomasi, C., and Manduchi, R. 1998. Bilateral filtering for gray and color images. In *ICCV*, 839–846. IEEE.
- Wang, Y.; Yang, J.; Yin, W.; and Zhang, Y. 2008. A new alternating minimization algorithm for total variation image reconstruction. *SIAM Journal on Imaging Sciences* 1(3):248–272.
- Wang, G.; Wong, T.-T.; and Heng, P.-A. 2006. Deringing cartoons by image analogies. *ToG* 25(4):1360–1379.
- Xu, L.; Lu, C.; Xu, Y.; and Jia, J. 2011. Image smoothing via l0 gradient minimization. In *ToG*, volume 30, 174. ACM.
- Xu, L.; Yan, Q.; Xia, Y.; and Jia, J. 2012. Structure extraction from texture via relative total variation. *ToG* 31(6):139.
- Xu, L.; Zheng, S.; and Jia, J. 2013. Unnatural l0 sparse representation for natural image deblurring. In *CVPR*, 1107–1114.
- Yang, J.; Ye, X.; Li, K.; Hou, C.; and Wang, Y. 2014. Color-guided depth recovery from rgb-d data using an adaptive autoregressive model. *TIP* 23(8):3443–3458.
- Zhang, Q.; Shen, X.; Xu, L.; and Jia, J. 2014. Rolling guidance filter. In *ECCV*, 815–830. Springer.
- Zhang, F.; Dai, L.; Xiang, S.; and Zhang, X. 2015. Segment graph based image filtering: fast structure-preserving smoothing. In *ICCV*, 361–369.
- Zhang, Z.; Kwok, J. T.; and Yeung, D.-Y. 2004. Surrogate maximization/minimization algorithms for adaboost and the logistic regression model. In *ICML*, 117. ACM.