Computer Science and Artificial Intelligence Laboratory

Technical Report

MIT-CSAIL-TR-2011-049 November 15, 2011

Fast and Robust Pyramid-based Image Processing

MATHIEU AUBRY, SYLVAIN PARIS, SAMUEL W. HASINOFF, JAN KAUTZ, and FR DO DURAND

massachusetts institute of technology, cambridge, ma 02139 usa — www.csail.mit.edu

Fast and Robust Pyramid-based Image Processing

MATHIEU AUBRY
INRIA / ENS
and
SYLVAIN PARIS
Adobe
and
SAMUEL W. HASINOFF
Google Inc.
and
JAN KAUTZ
University College London
and
FRÉDO DURAND
Massachusetts Institute of Technology

Multi-scale manipulations are central to image editing but they are also prone to halos. Achieving artifact-free results requires sophisticated edge-aware techniques and careful parameter tuning. These shortcomings were recently addressed by the local Laplacian filters, which can achieve a broad range of effects using standard Laplacian pyramids. However, these filters are slow to evaluate and their relationship to other approaches is unclear. In this paper, we show that they are closely related to anisotropic diffusion and to bilateral filtering. Our study also leads to a variant of the bilateral filter that produces cleaner edges while retaining its speed. Building upon this result, we describe an acceleration scheme for local Laplacian filters that yields speed-ups on the order of 50×. Finally, we demonstrate how to use local Laplacian filters to alter the distribution of gradients in an image. We illustrate this property with a robust algorithm for photographic style transfer.

Categories and Subject Descriptors:

General Terms: Photo editing, Image processing

Additional Key Words and Phrases: Computational photography, photo editing, image processing, Laplacian pyramid, bilateral filter, photographic style transfer

ACM Reference Format:

1. INTRODUCTION

Manipulating images at multiple scales is a challenging task. Direct linear manipulation of the frequency bands yields unsightly halos. This issue has been addressed by many nonlinear approaches. While these techniques produce halo-free images, they also come with their own shortcomings such as limited scalability due to the need of solving a global optimization problem, e.g. [Fattal et al. 2002; Farbman et al. 2008; Subr et al. 2009], or edge defects that require corrections in post-process [Durand and Dorsey 2002; Bae et al. 2006; Kass and Solomon 2010]. Recently, Paris et al. [2011] described the *local Laplacian filters* that address these shortcomings and produce high-quality results over a wide range of parameters. However, while these filters achieve similar effects to existing edge-aware filters, their relationship to other approaches is unclear, leaving crucial questions such as "why do they work?" and "what do they do differently?" largely unanswered. Further, these filters are prohibitively slow in their original form. Paris and colleagues [2011] mitigate this issue with a heuristic approximation but its properties and accuracy are unknown, and even so, it remains slow.

In this paper, we study these filters to gain a better understanding of their behavior. First, we rewrite them as the averaging at each scale of the signal variations in the local neighborhood around each pixel. From this formulation, we show that local Laplacian filters can be interpreted as a multi-scale version of anisotropic diffusion, and that they are closely related to bilateral filtering, the main difference being how they are normalized. While the difference is minor in uniform regions, it becomes large in configurations such as edges, corners, and isolated pixels where the bilateral filter is known not to perform well [Durand and Dorsey 2002; Buades et al. 2006]. We use this insight to design a variant of the bilateral filter, which we name *unnormalized bilateral filtering*, and we show that it yields significantly cleaner edges. We also propose a signal-processing interpretation of local Laplacian filtering applied to gray-scale images and derive a new acceleration scheme grounded on sampling theory. Our analysis shows that we can quantize the intensity scale while introducing only negligible differences with the results produced by the original scheme. Our tests show that our algorithm is about 50 times faster than the heuristics of Paris et al. and that it runs at interactive rates on CPUs without resorting to parallelism. Further, our GPU implementation processes one-megapixel images at 20 Hz. Finally, we show how to use these filters to alter the gradient distribution of an image without introducing halos. We illustrate this approach in the context of photographic style transfer (Fig. 1). Our experiments show that our method achieves satisfying transfers on a larger set of photos than previous work.

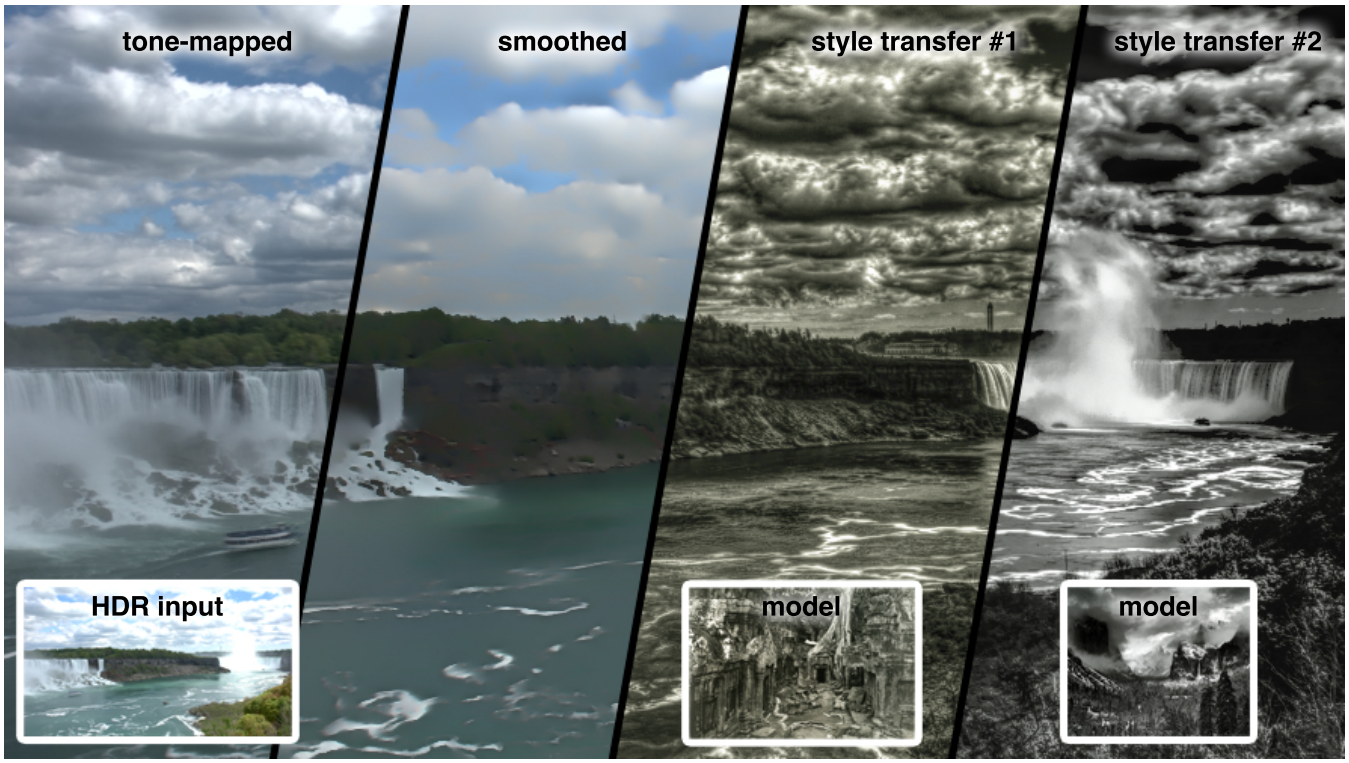

Fig. 1: *In this paper, we show that the local Laplacian filter [Paris et al. 2011] is related to anisotropic diffusion and that it can be also understood as a multiscale variant of bilateral filtering. This enables us to derive an efficient algorithm that is* $50\times$ *faster, achieving, for instance, interactive tone-mapping (left) and detail reduction (second left). Our analysis also leads to a better understanding of local Laplacian filtering, which allows us to modify gradient distributions without artifacts, yielding robust and efficient style transfer (two examples on right, models inset at the bottom).*

We make the following contributions.

—We formally characterize the similarities between local Laplacian filters, anisotropic diffusion and bilateral filtering.

—Based on this understanding, we build a new single-scale filter, the unnormalized bilateral filter, that behaves similarly to the bilateral filter in smooth areas while producing cleaner edges.

—We describe a fast algorithm for Laplacian filtering on gray-scale images that is about 50 times faster than the heuristic of Paris et al. and that is guaranteed to approximate the original scheme.

—We explain how to transfer the gradient histogram from one image to another using Laplacian filtering and demonstrate its application to photographic style transfer.

1.1 Related work

Relationship between nonlinear filters. A dense net of studies now relate methods as diverse as bilateral filtering, anisotropic diffusion, mean shift, neighborhood filtering, mode filtering, and robust statistics, e.g. [Black et al. 1998; Durand and Dorsey 2002; Elad 2002; van de Weijer and van den Boomgaard 2002; Barash and Comaniciu 2004; Buades et al. 2005; Mrzek et al. 2006; Paris et al. 2009]. However, the recently proposed local Laplacian filters are not yet part of this mesh and little is known about their relationship to the existing body of work on nonlinear image filtering. A contribution of our paper is to show that they are closely related to anisotropic diffusion and to the bilateral filter.

Bilateral filtering. The bilateral filter [Tomasi and Manduchi 1998] is a popular edge-aware smoothing filter for computational photography applications, e.g. [Durand and Dorsey 2002; Bae et al. 2006; Paris et al. 2009], because it achieves satisfying results while being fast [Chen et al. 2007; Paris and Durand 2009; Adams et al. 2009; Adams et al. 2010]. However, it is also known to suffer from over-sharpening, which introduces unsightly edge defects [Buades et al. 2006] and requires applying a fix in post-processing, e.g. [Durand and Dorsey 2002; Bae et al. 2006; Kass and Solomon 2010]. This additional step requires more computation and introduces new parameters to set. In comparison, our variant modifies the original bilateral filter in minor ways that preserve its speed and ease of use, while significantly reducing over-sharpening.

Local Laplacian filtering. Paris et al. [2011] introduced local Laplacian filtering as an alternative to existing edge-aware filters. While they demonstrate high-quality results, the running times are slow, on the order of a minute per megapixel with a single thread, which requires a parallel implementation and an approximation scheme to reach interactive rates. Whereas the effects of this approximation are unclear, we propose an acceleration technique firmly grounded on signal processing analysis [Chen et al. 2007; Paris and Durand 2009], which allows us to control the trade-off between speed and accuracy. And as previously discussed, we also describe a theoretical relationship between local Laplacian filters, anisotropic diffusion, and bilateral filtering. Since local Laplacian

Variable	Description	Variable	Description
$\mathbf{p} = (x, y)$	spatial location	$L_\ell[I]$	level ℓ of pyramid
ℓ	level in pyramid	$\{L_\ell[I]\}$	Laplacian pyramid
I	input image	g_0	Gaussian coefficient
O	output image	G_σ	Gaussian kernel
$r(i)$	remapping func.		

Table I. : *Common notation used in this paper.*

filters are at the core of our work, we describe them in more detail in Section 1.2.

Photographic style transfer. Bae et al. [2006] transfer the "look" of one photographer's masterpiece onto another photo by matching statistics such as the intensity and texture histograms of the two pictures. While they demonstrate convincing results, the method consists of many steps, including solving the Poisson equation several times, which limits the ability of the approach to process high-resolution images and makes the technique difficult to implement and tune. Sunkavalli et al. [2010] propose a simpler alternative based on image pyramids but their results do not match the look of the model photograph as well. In this paper, we demonstrate that Laplacian filtering can be used for manipulating the gradient histogram of an image. Our approach generates visual matches in the same spirit as Bae's technique. For strongly stylized examples, it often performs better, because the robustness of local Laplacian filters allows for larger image modifications.

1.2 Background on local Laplacian filters

Local Laplacian filters are edge-aware operators that define the output image O by constructing its Laplacian pyramid $\{L_\ell[O]\}$ coefficient by coefficient. To compute a given coefficient $L_{\ell_0}[O](x_0, y_0)$, one first processes the input image I with a point-wise nonlinearity $r(\cdot)$ that depends on the local properties of I around (x_0, y_0). Intuitively, $r(I)$ is an image that looks like the desired result but only locally at position (x_0, y_0) and scale ℓ_0. For instance, to increase the amount of detail, the local contrast around (x_0, y_0) is boosted with a local S-shaped tone curve centered on the value $g_0 = G_{\ell_0}[I](x_0, y_0)$ where $G_{\ell_0}[I]$ is the ℓ_0 level of the Gaussian pyramid of I. This remapping increases the local intensity variations while preserving the local intensity average (Fig. 2), i.e. g_0 is not altered by the curve. This may compress the contrast at other locations but local Laplacian filtering only needs $r(I)$ to look good locally. We will discuss the formal definition of r later. Given $r(I)$ for a particular position (x_0, y_0) and scale ℓ_0, the pyramid coefficient in the output $L_{\ell_0}[O](x_0, y_0)$ is defined as the corresponding coefficient in the remapped input image $L_{\ell_0}[r(I)](x_0, y_0)$.

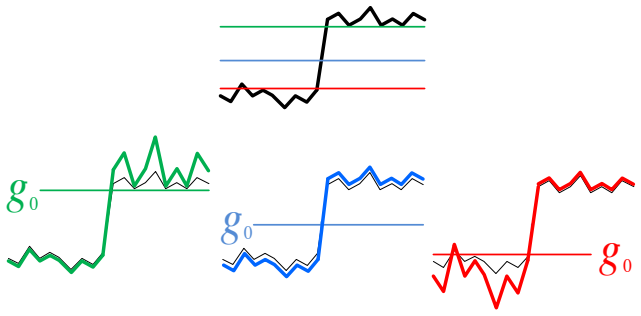

Fig. 2: *Effect of a detail enhancement ($0 \le \alpha < 1$) remapping function r with several reference values g_0 in the neighborhood of an edge. Details are enhanced for values similar to the reference value, but not for values far from it.*

A direct application of this approach yields an $\mathcal{O}(N^2)$ algorithm where N is the number of image pixels, but only processing the portion of I needed to compute $L_{\ell_0}[r(I)](x_0, y_0)$ reduces the complexity to $\mathcal{O}(N \log N)$. Paris et al. [2011] further accelerate the process and obtain an $\mathcal{O}(N)$ method by using heuristic that amounts to applying r to a downsampled version of I when processing coarse pyramid levels. While satisfying results are obtained in practice, there is no clear understanding of the speed-versus-accuracy trade-off inherent in this heuristic and a multi-core implementation is still required to achieve interactive running times. A contribution of our work is to propose a novel acceleration scheme that is faster and grounded on signal processing theory. For the nonlinearity r, the original article focuses only on a small set of options defined by three parameters. While these are sufficient for detail manipulation and tone mapping, in this paper, we reinterpret r in terms of first-order image statistics and explore more general designs in the context of style transfer.

Design of the remapping function. To compute the coefficient (ℓ_0, x_0, y_0), Paris and colleagues propose the following remapping functions:

$$r(i) = \begin{cases} g_0 + \text{sign}(i - g_0)\, \sigma_r\, (|i - g_0| / \sigma_r)^\alpha & \text{if } i \le \sigma_r \\ g_0 + \text{sign}(i - g_0) \left(\beta\, (|i - g_0| - \sigma_r) + \sigma_r\right) & \text{if } i > \sigma_r \end{cases} \quad (1)$$

where g_0 is the coefficient of the Gaussian pyramid at (ℓ_0, x_0, y_0), which acts as a reference value, α controls the amount of detail increase ($0 \le \alpha < 1$) or decrease ($\alpha > 1$), β controls the dynamic range compression ($0 \le \beta < 1$) or expansion ($\beta > 1$), and σ_r defines the intensity threshold the separates details from edges. Sample functions are plotted in Figure 3. Another contribution of our paper is to describe a new class of remapping functions that enable gradient histogram transfer.

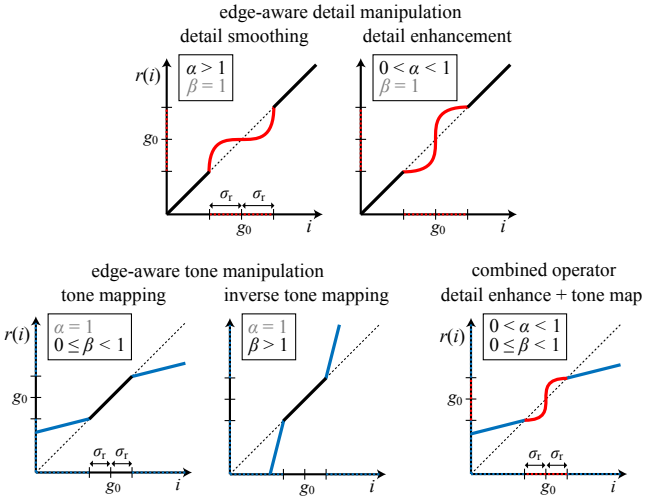

Fig. 3: *Remapping functions proposed by Paris et al. [2011]. Reproduced from [Paris et al. 2011]*

2. BILATERAL FILTERING, ANISOTROPIC DIFFUSION, AND LOCAL LAPLACIAN FILTERS

In this section, we study the local Laplacian filter and relate it to anisotropic diffusion and bilateral filtering. We start by formalizing this relationship and then define a new variant of the bilateral filter inspired by this result.

Background. With our notation, anisotropic diffusion as introduced by Perona and Malik [1990] is defined by a partial differential equation:

$$\frac{\partial I}{\partial t} = \operatorname{div}\big(w(\nabla I)\,\nabla I\big) \tag{2}$$

where t represents the diffusion time, i.e. how long the process has been run, and w is a weighting function that is equal to 1 for $\nabla I = \mathbf{0}$ and decreases for larger gradients. This equation is discretized as an iterative process:

$$I_{t+1}(\mathbf{p}) \;=\; I_t(\mathbf{p}) \;+\; \sum_{\mathbf{q}\in\mathcal{N}_4(\mathbf{p})} w\big(I_t(\mathbf{q})-I_t(\mathbf{p})\big)\,\big[I_t(\mathbf{q})-I_t(\mathbf{p})\big] \tag{3}$$

where t now counts how many iterations have been performed, and $\mathcal{N}_4$ is the 4 neighborhood of $\mathbf{p}$. The process is initialized at $t=0$ with I_0 being the input image.

Using $I_\mathbf{p}$ as a shorthand for $I(\mathbf{p})$, the bilateral filter is defined as [Tomasi and Manduchi 1998]:

$$BF_\mathbf{p} = \frac{1}{W_\mathbf{p}} \sum_\mathbf{q} G_{\sigma_\mathrm{s}}(\mathbf{q}-\mathbf{p})\, G_{\sigma_\mathrm{r}}(I_\mathbf{q}-I_\mathbf{p})\, I_\mathbf{q} \tag{4a}$$

$$W_\mathbf{p} = \sum_\mathbf{q} G_{\sigma_\mathrm{s}}(\mathbf{q}-\mathbf{p})\, G_{\sigma_\mathrm{r}}(I_\mathbf{q}-I_\mathbf{p})\,, \tag{4b}$$

where G_{σ_r} and G_{σ_s} are Gaussian kernels of variance σ_r^2 and σ_s^2 defined by $G_\sigma(x) = \exp(-x^2/2\sigma^2)$ that are called the *range weight* and *space weight* respectively. Formally, the sums $\sum_\mathbf{q}$ cover the entire image but in practice are limited to local windows of radius $3\sigma_\mathrm{s}$ since G_{σ_s} becomes almost zero for distant pixels. Using the symmetry of the Gaussian kernel and the fact that the weights sum up to 1, Equation 4a can be rewritten to make bilateral filtering appear as a multi-scale version of anisotropic diffusion as described in Equation 2 with $w = G_{\sigma_\mathrm{r}}$ [Elad 2002; Barash and Comaniciu 2004]:

$$BF_\mathbf{p} = I_\mathbf{p} + \frac{1}{W_\mathbf{p}} \sum_{d>0} G_{\sigma_\mathrm{s}}(d) \sum_{\substack{\mathbf{q}\ \text{s.t.}\\ \|\mathbf{q}-\mathbf{p}\|=d}} G_{\sigma_\mathrm{r}}(I_\mathbf{q}-I_\mathbf{p})\,(I_\mathbf{q}-I_\mathbf{p}) \tag{5}$$

As we shall see later, it is also convenient to rewrite Equation 4a as:

$$BF_\mathbf{p} = I_\mathbf{p} + \frac{1}{W_\mathbf{p}} \sum_\mathbf{q} G_{\sigma_\mathrm{s}}(\mathbf{q}-\mathbf{p})\, G_{\sigma_\mathrm{r}}(I_\mathbf{q}-I_\mathbf{p})\,(I_\mathbf{q}-I_\mathbf{p})\,. \tag{6}$$

2.1 Single scale local Laplacian filter

In this section, we relate local Laplacian filtering to anisotropic diffusion and bilateral filtering. We show that the Gaussian kernel used to build the image pyramids acts as a spatial weight and that the remapping function r as a range weight, thereby resembling bilateral filtering. We can further decompose the spatial kernel into rings of fixed radius to make appear a link with anisotropic diffusion akin to the studies by Elad [2002] and Barash and Comaniciu [2004].

For this study, we consider the space of remapping functions with the form

$$r(i) = i - (i-g_0)\, f(i-g_0)\,, \tag{7}$$

where f is a continuous function. This space includes the functions of Paris et al. (Eq. 1) as a special case. We first consider a two-level pyramid, that is, we seek to compute the levels $L_0[O]$ and $L_1[O]$ of the Laplacian pyramid of the output image O. We assume for now that the residual remains unprocessed, that is $L_1[O] = L_1[I]$. For a pixel $\mathbf{p} = (x_0, y_0)$ on the 0th level we have:

$$L_0[O](\mathbf{p}) = r(I_\mathbf{p}) - \big[\bar{G}_{\sigma_\mathrm{p}} * r(I)\big](\mathbf{p})\,, \tag{8}$$

where $\bar{G}_{\sigma_\mathrm{p}} = \frac{1}{\sqrt{2\pi\sigma_\mathrm{p}^2}} G_{\sigma_\mathrm{p}}$ is a normalized Gaussian kernel of variance σ_p^2 used to build the pyramids, and $*$ is the convolution operator. Expanding r and using $L_0[I] = I - \bar{G}_{\sigma_\mathrm{p}} * I$ and $g_0 = I_\mathbf{p}$ since we are at the finest level of the pyramid, we obtain:

$$L_0[O](\mathbf{p}) = L_0[I](\mathbf{p}) + \big[\bar{G}_{\sigma_\mathrm{p}} * (I-I_\mathbf{p})\, f(I-I_\mathbf{p})\big](\mathbf{p})\,. \tag{9}$$

Then, by upsampling the unmodified residual, adding it to both sides, and expanding the convolution, we get the formula that we seek:

$$O_\mathbf{p} = I_\mathbf{p} + \sum_\mathbf{q} \bar{G}_{\sigma_\mathrm{p}}(\mathbf{q}-\mathbf{p})\, f(I_\mathbf{q}-I_\mathbf{p})\,(I_\mathbf{q}-I_\mathbf{p}) \tag{10}$$

This equation shows that the local Laplacian filter computes a local average in the same spirit as the bilateral filter (Eq. 6) using $\bar{G}_{\sigma_\mathrm{p}}$ as the spatial weight and f as the range weight. If we choose $f = G_{\sigma_\mathrm{r}}$, then the two-level local Laplacian filter becomes almost the same as the bilateral filter—the only difference is that the weights are *not* normalized by $\frac{1}{W_\mathbf{p}}$. This simple modification defines a new filter we call the *unnormalized bilateral filter* which we examine in Section 2.2.

Equation 10 can also be written as follows using the symmetry of the Gaussian kernel [Elad 2002; Barash and Comaniciu 2004]:

$$O_\mathbf{p} = I_\mathbf{p} + \sum_{d>0} \bar{G}_{\sigma_\mathrm{p}}(d) \sum_{\substack{\mathbf{q}\ \text{s.t.}\\ \|\mathbf{q}-\mathbf{p}\|=d}} f(I_\mathbf{q}-I_\mathbf{p})\,(I_\mathbf{q}-I_\mathbf{p}) \tag{11}$$

This formulation shows that, similar to the bilateral filter (Eq. 6), a two-scale local Laplacian filter can be seen as a multiscale version of anisotropic diffusion as described in Equation 2 [Perona and Malik 1990]. The main difference between bilateral filtering and Laplacian filtering is how each scale is weighted. The bilateral filter uses weights that sum to 1 because of the $1/W_\mathbf{p}$ normalization factor whereas the local Laplacian filters apply unnormalized weights that do not sum to 1.

In the case of more than two levels, g_0 is not equal to $I_\mathbf{p}$, and we cannot collapse the pyramid as above. Nonetheless, we can write:

$$L_{\ell_0}[O] = D_{\ell_0} * (I-g_0)\, f(I-g_0)\,, \tag{12}$$

where $D_\ell = \bar{G}_{2^{\ell-1}\sigma_\mathrm{p}} - \bar{G}_{2^\ell\sigma_\mathrm{p}}$ is the difference-of-Gaussians filter defining the pyramid coefficients at level ℓ. This expression can be rewritten as

$$L_{\ell_0}[O](\mathbf{p}) = \sum_\mathbf{q} D_{\ell_0}(\mathbf{q}-\mathbf{p})\, f(I_\mathbf{q}-g_0)\,(I_\mathbf{q}-g_0)\,. \tag{13}$$

This shows that each level of the output pyramid is a local average of differences over a neighborhood of $\mathbf{p}$. This result will be useful in our design of an acceleration scheme in Section 3.

2.2 Unnormalized bilateral filtering

In the previous analysis, we showed that local Laplacian filters share similarities with bilateral filtering. But they are not identical; Laplacian filters are multiscale and do not normalize the contribution of the pixels, whereas the bilateral filter is two-scale and

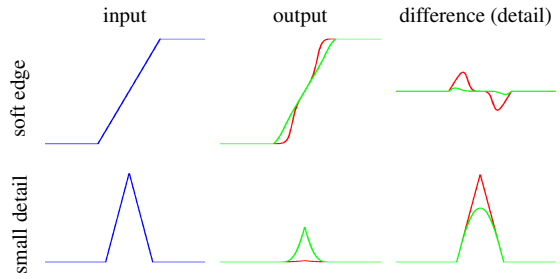

Fig. 4: *We compare how the standard bilateral filter (red) and our unnormalized variant (green) process simple features. The top row shows a soft edge (left). The standard filter over-sharpens the edge (center) and introduces non-negligible content in the detail layer (right). These spurious variations are a source of halos and gradient reversals in applications such as tone mapping. In comparison, the unnormalized filter produces a cleaner edge with almost no spurious content in the detail layer. The bottom row shows an isolated small detail (left), with a different vertical scale, whose height is 10% of the edge amplitude. The standard bilateral filter smooths such details more aggressively than the unnormalized version. This explains why the tone-mapping results of the unnormalized version tend to be softer, but this is a minor side effect compared to the creation of edge artifacts. Both filters were set with σ_r equal to 10% of the edge amplitude.*

normalized. This suggests a variant of the bilateral filter where we remove the overall normalization but keep the two-scale design. We call it the *unnormalized bilateral filter*. Formally, we define it as:

$$UBF_{\mathbf{p}} = I_{\mathbf{p}} + \sum_{\mathbf{q}} \bar{G}_{\sigma_s}(\mathbf{q}-\mathbf{p})\, G_{\sigma_r}(I_{\mathbf{q}} - I_{\mathbf{p}})\, (I_{\mathbf{q}} - I_{\mathbf{p}}) \qquad (14)$$

Compared to the bilateral filter, this unnormalized version has a weaker effect when the sum of the weights W is small. This occurs when the center pixel is different from many of its neighbors, which typically happens at edges. Durand and Dorsey [2002] interpret W from a robust statistics standpoint and explain that when it is small, the bilateral filter returns an estimate based on limited data, which causes the artifacts that appear at some strong edges. The unnormalized version can be interpreted as a filter that is weaker at those ambiguous locations (Fig. 4). This tends to generate slightly softer images but greatly reduces the artifacts as shown in Figure 5 and in supplemental material. We also compared to the post-process fix described by Durand and Dorsey that blends the bilateral filter output with a blurred version of the input using $\log(W)$ as blending control. Similarly, the unnormalized filter achieves cleaner and softer outputs without requiring post-processing. Although the results are not as detailed as with the multiscale local Laplacian filter, the unnormalized filter is about 5 times faster, which can make it a useful alternative.

Acceleration method. We can compute the unnormalized bilateral filter even more efficiently using an acceleration scheme inspired by the bilateral grid [Chen et al. 2007; Paris and Durand 2009]. This scheme computes expressions of the form $\sum_{\mathbf{q}} w_s(\mathbf{q}-\mathbf{p})\, w_r(I_{\mathbf{q}} - I_{\mathbf{p}})\, X$ at discrete grid points, where w_s and w_r are space and range weights respectively, and X represents data defined over the image domain. The unnormalized bilateral filter corresponds to $w_s(x) = \bar{G}_{\sigma_s}(x)$, $w_r(x) = x G_{\sigma_r}(x)$, and $X = 1$ everywhere. The bilateral grid yields accurate results as long as w_r and w_s are band-limited, which is the case with the unnormalized bilateral filter. While the standard bilateral filter requires two grids, one for the intensity average and one for the normalization factor, the unnormalized version needs only one grid since there is no normalization factor, which makes it twice as fast.

3. EFFICIENT LOCAL LAPLACIAN FILTERING

We propose an acceleration technique to evaluate local Laplacian filters on single-channel images. This encompasses many practical cases such as detail manipulation and tone mapping [Paris et al. 2011] as well as photographic style transfer that we discuss later in Section 4. Our strategy is based on the fact that the nonlinearity comes from the dependency on g_0. We characterize this dependency in terms of signal processing, which allows us to design a theoretically grounded subsampling scheme that is more than an order of magnitude faster than the algorithm proposed by Paris et al. [2011]. In practice, we precompute a small set of pyramids $\{L_\ell[r_j(I)]\}$ over different values γ_j of g_0, where r_j is the remapping function for $g_0 = \gamma_j$. Whenever we need a pyramid coefficient for a particular g_0 value, instead of remapping the image and computing a new pyramid, which is expensive, we find j such that $\gamma_j \le g_0 < \gamma_{j+1}$ and interpolate the coefficients of precomputed pyramids j and $j+1$. Formally, we seek to sample r as sparsely as possible without losing accuracy. If r is band-limited, using the sampling theorem, the optimal sampling is the Nyquist limit, i.e.,

(a) HDR input (b) uncorrected BF (c) corrected BF (d) unnormalized BF (e) local Laplacian filter

Fig. 5: *Compressing an HDR image (a) with strong detail enhancement with the bilateral filter (b) leads to artifacts on the border of strong edges, e.g. on the window structure and on the lamp base. Some of them can be fixed by a postprocessing step (c), for instance the window structure is significantly improved–other parts such as the lamp base remains problematic. With our unnormalized bilateral filer (d), most these artifacts are avoided although the rendition is overall softer, and none appear with our fast version of the local Laplacian filter (e). However, the local Laplacian filter remains slower than the unnormalized filter. (b,c,d,e) are close-ups of the lamp at the end of corridor (a).*

half the smallest wavelength present in the signal. To estimate this value, we observe r as a function of g_0. From that perspective, only the term $(i-g_0)\,f(i-g_0)$ is not constant and what actually matters is the frequency content of $x\,f(x)$. Denoting the Fourier transform by $\mathcal{F}[\cdot]$ and using $'$ for derivatives, the property $\mathcal{F}[x\,f(x)] \propto \mathcal{F}[f]'$ ensures that, if f is band-limited, r is as well. This means that, if f is band-limited, we can sparsely sample the intensity domain with only minimal loss. We further discuss the accuracy of this approach at the end of this section.

Our algorithm is as follows:

(1) Compute the Gaussian pyramid of I.

(2) Regularly sample the intensity range with the $\{\gamma_j\}$ values.

(3) Compute the remapped images $\{r_j(I)\}$ and their Laplacian pyramids $\{L_\ell[r_j(I)]\}$.

(4) For each pyramid coefficient (ℓ_0, x_0, y_0):
   i. Get the corresponding coefficient g_0 in the Gaussian pyramid.
   ii. Compute a and j such that $g_0 = (1-a)\gamma_j + a\,\gamma_{j+1}$.
   iii. Linearly interpolate the output coefficient from the precomputed pyramids: $L_{\ell_0}[O](x_0,y_0) = (1-a)L_{\ell_0}[r_j(I)](x_0,y_0) + aL_{\ell_0}[r_{j+1}(I)](x_0,y_0)$.

(5) Collapse the output pyramid $\{L_\ell[O]\}$.

In practice, we use a Gaussian function G_σ for f and we recommend sampling the intensity range every standard deviation σ.

Because the number of precomputed pyramids is fixed, this algorithm has linear complexity in the number of pixels. Since our implementation uses a Gaussian for f, r is not strictly speaking band-limited. Further, we use linear interpolation instead of a sinc kernel for reconstructing the signal. The effect of these simplifications is that our algorithm does not perfectly reproduce the result of the original algorithm. Nevertheless, it produces accurate approximations above 30 dB and the differences are invisible in practice. Most importantly, when sampling every standard deviation, it runs at interactive rates, about 350 ms per megapixel on a 2.66 GHz Intel Core i7, which is about 50× faster than the heuristic of Paris and colleagues for the same accuracy (Fig. 6). We also ported our algorithm to graphics hardware (an NVIDIA GeForce 480 GTX), where it runs at about 49 ms for a one megapixel image and about 116 ms for a four megapixel image, which is about 10 times faster than Paris's heuristic implemented on the same card.

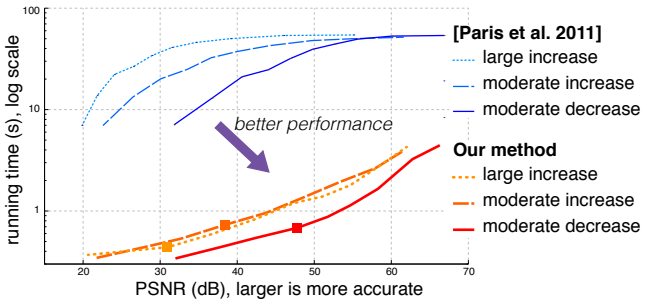

Fig. 6: *We measured the running time and accuracy of Paris's acceleration scheme and ours for several settings. For Paris's algorithm, we varied the depth of the intermediate sub-pyramids; for ours, we varied the sampling rate of the intensity range. For each setting of both methods, we applied three different sets of parameters to achieve a large detail increase, a moderate increase, and a moderate decrease. For the same accuracy, our scheme is about* 50× *faster than Paris's. The square marks on the curve of our method indicate the one-sample-per-standard-deviation sampling rate. We used a* 1600×1200 *image.*

4. PHOTOGRAPHIC STYLE TRANSFER

As we showed in Section 2, the re-centering of the remapping makes local Laplacian filters rely on differences between the value of a pixel and its neighbors. This makes them closely related to image gradients. In this section, we propose a new method using local Laplacian filters to alter the *distribution* of these differences and, by extension, the distribution of image gradients. Although our approach is mostly empirical, it builds upon a strong intuition and performed well in practice. We illustrate this capability with an algorithm to transfer photographic style akin to the method of Bae et al. [2006].

4.1 Manipulating gradient distributions

To gain intuition, we first consider a two-level pyramid and further assume that the pixel $\mathbf{p}$ has a single neighbor $\mathbf{q}$ and ignore the spatial weight $\tilde{G}_{\sigma_p}$. The output of the filter then becomes: $O_\mathbf{p} = I_\mathbf{p} + f(I_\mathbf{q} - I_\mathbf{p})\,(I_\mathbf{q} - I_\mathbf{p})$. We highlight the role of pixel differences by subtracting $I_\mathbf{q}$ on both sides. Assuming that f is symmetric, which is always the case in practice, we get: $O_\mathbf{p} - I_\mathbf{q} = (I_\mathbf{p} - I_\mathbf{q}) - f(I_\mathbf{p} - I_\mathbf{q})\,(I_\mathbf{p} - I_\mathbf{q})$. Defining $h(x) = [1 - f(x)]x$, this can be rewritten in the more concise form $O_\mathbf{p} - I_\mathbf{q} = h(I_\mathbf{p} - I_\mathbf{q})$ which shows that the filter remaps $I_\mathbf{p}$ so that its difference with its neighbor $I_\mathbf{q}$ has a desired value specified by the h function. Since $\mathbf{p}$ and $\mathbf{q}$ are neighbors, this can be seen as remapping the image gradient at $\mathbf{p}$. If we now consider again a larger neighborhood as in Equation 10, the filter can be interpreted as making a trade-off between the desired gradient values coming from different neighbors. $\tilde{G}_{\sigma_p}$ weights the contribution of each pixel $\mathbf{q}$ and h defines the desired output gradients. Further, h is sufficient to define a local Laplacian filter since $r(i) = g_0 + h(i - g_0)$ where r is the remapping function originally defined by Paris et al. (§ 1.2). This comes directly from the definitions of f and h.

Building upon this intuition, we describe a method to transfer the histogram of gradient amplitudes from a model image M to the input image I. We apply local Laplacian filtering with a remapping function r defined such that gradient statistics of M are transferred to I. For both images, we compute the histogram of the gradient amplitudes $\|\nabla I\|$ and $\|\nabla M\|$ and the corresponding histogram transfer function t, i.e. $t(x) = CDF\big[\|\nabla M\|\big]^{-1}\big(CDF\big[\|\nabla I\|\big](x)\big)$. This function t means that when two pixels have a difference with amplitude $|x|$ in I, we want a difference with amplitude $t(|x|)$ in O. Finally, we seek to preserve the sign of the difference, and we define $h(x) = \text{sign}(x)\,t(|x|)$ that leads to the remapping function:

$$r(i) = g_0 + \text{sign}(i - g_0)\,t(|i - g_0|) \tag{15}$$

We use this function to run a local Laplacian filter on I. While we built our intuition on a simplified case based only on two levels and two pixels, in practice, the situation is more complex and we need to iterate to obtain the desired result. Figures 7 and 8a show that our approach quickly converges after a few iterations. The only fixed point of the iteration is when r is the identity function, which implies that t is also the identity function, which finally implies $CDF\big[\|\nabla M\|\big] = CDF\big[\|\nabla I\|\big]$. This guarantees that the process can only converge to the desired result where the output has the gradient histogram of the model.

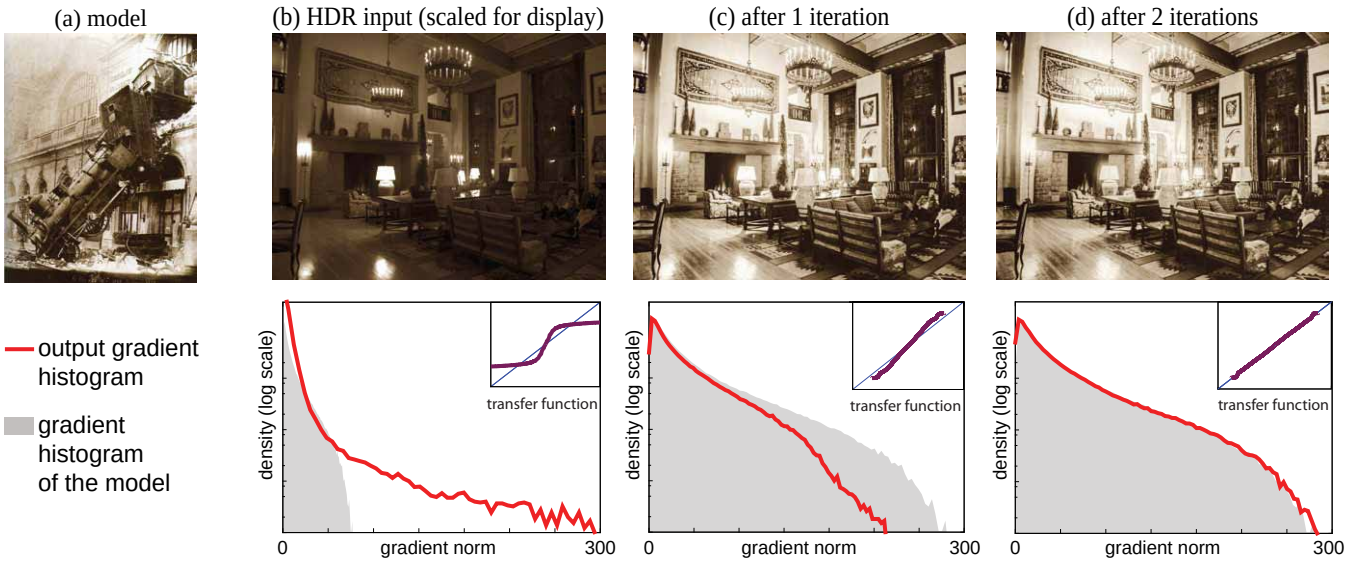

Fig. 7: *Our iterative method to transfer gradient histograms stabilizes quickly. Visually, the results do not change after 2 iterations.*

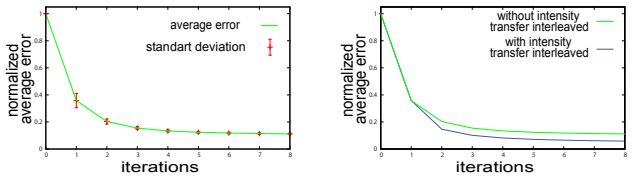

Fig. 8: *To validate our histogram transfer method, we measured the Earth Mover's distance between the gradient histogram of the output image and the gradient histogram of the model image as a function of the number of iterations. We report the average and standard deviation of this distance over 15 image pairs. We normalized the distances so that the difference between the input and model histograms is 1. The output histogram quickly becomes closer to the model and then the convergence slows down (a). We tested more iterations, up to 1000, the distance keeps reducing but slowly. Visually the results becomes stable after a few iterations (Fig. 7). Interleaving standard intensity histogram transfer with our local Laplacian method speeds up the convergence at a minor computational cost (b).*

4.2 Style transfer algorithm

We demonstrate how to achieve photographic style transfer in the spirit of Bae et al. [2006] using the gradient transfer method described in the previous section. We seek to transfer the "look" of the model image M to an input image I. Typically, M is a picture by a master, such as Ansel Adams, and I is a photo by a casual photographer who wishes to mimic the master's style. Bae's technique involves solving the Poisson equation twice to mitigate the over-sharpening artifacts inherent in bilateral filtering. This has two drawbacks: first, the global optimization limits scalability, and second, these corrections can limit the large image transforms necessary to achieve more extreme looks. In comparison, our method is optimization-free and its increased robustness enables more strongly stylized renditions.

Our algorithm follows the same overall approach as Bae's: we seek to match both the global contrast, i.e. the large-scale intensity variations, and the local contrast, i.e. the amount of texture of the model image M. We use an iterative process. For each iteration, we first compute the histograms of the input and model gradients, $\|\nabla I\|$ and $\|\nabla M\|$. We build the transfer remapping function r as described in Section 4.1 and apply the corresponding local Laplacian filter. Then, we apply a standard intensity histogram transfer to match the intensity distribution of the model M. We typically apply a few such iterations, 4 in all the results presented in this paper and in supplemental material. Intuitively, the local Laplacian filtering step transfers the local contrast, i.e. the gradients, and the intensity matching step transfer the global variations, i.e. the intensity distribution. Further, we found empirically that interleaving the histogram matching step speeds up the convergence of the gradient transfer as shown in Figure 8b. We end by a histogram matching step. This favors matching global contrast or local contrast, which we found visually more satisfying. More importantly, gradient transfer alone can produce out-of-range values resulting in over-exposed regions (Fig. 7). Ending with intensity matching ensures that the output image has the same dynamic range as the model, thereby preventing overexposure.

4.3 Results

We demonstrate the robustness of our method using a variety of input and model images for which we successfully transfer the look of the artist's photo to the input (Fig. 11 and 12). Our technique handles standard images and HDR images seamlessly since it produces output with the same dynamic range as the model photograph.

A limitation of our approach is the lack of semantic understanding of the scene. For instance, the method can sometimes introduce unnatural variations in a uniform sky or darken regions that one would expect to be lit (Fig. 9). Fixing these problems in general requires either the user in the loop or a higher-level analysis of the scene, both of which are beyond the scope of this paper but would be interesting future work. That said, these cases are rare and our

Fig. 9: *Failure cases. Our algorithm only has a low-level view of an image and does not know its semantics. In rare cases, it can produce unexpected results such as low-frequency variations in the sky (left) or darkening of a tower that we expect to be well lit (right).*

approach performs well most of the time. Compared to the method by Bae et al. [2006], our approach often performs better, especially in its ability to reproduce the style of texture. Because Bae and colleagues use the bilateral filter, they need to correct their results using strong gradient-domain constraints that limit their ability to modify the input image. This is visible when comparing Bae's results with and without imposing these constraints (Fig. 10). While Bae's results without constraints better match the models, they suffer from halos at strong edges. The gradient constraints mitigate these artifacts but come at the cost of significantly duller renditions. Since our approach does not require such strong constraints, we are able to obtain high-quality results that better reproduce the texture in the models without introducing halos. We also compared our method to histogram transfer applied to the gradient amplitudes followed by a Poisson reconstruction. As shown in Figure 11, this naïve approach matches the amount of details in the models poorly and does not yield satisfying results. Finally, we also experimented with the multiscale approach of Sunkavalli et al. [2010] and found that the produced image does not match the model look as well as our approach (Fig. 11 and 12).

(a) HDR input naive gamma compression (b) Bae et al. without gradient correction (c) Bae et al. with gradient correction (d) our method

Fig. 10: *Since the method by Bae et al. [2006] relies on the bilateral filter, it suffers from rim halos at strong edges (b). This can be corrected by manipulating the gradient field of the result. However, this slows down the computation and yields duller results (c). In comparison, our approach directly produces satisfying results and does not require solving a costly Poisson equation (d).*

5. CONCLUSION

We have studied local Laplacian filters and have shown that they are closely related to bilateral filtering and anisotropic diffusion. This insight has led to several practical contributions: we have described the unnormalized bilateral filter which improves the results of the bilateral filter at edges, sped up the local Laplacian filters, and described a technique to manipulate image gradients that leads to a robust algorithm for transferring photographic style. We believe that these improvements make local Laplacian filtering usable and suitable for interactive image editing.

ACKNOWLEDGMENTS

REFERENCES

ADAMS, A., BAEK, J., AND DAVIS, A. 2010. Fast high-dimensional filtering using the permutohedral lattice. *Computer Graphics Forum (Proc. Eurographics)*.

ADAMS, A., GELFAND, N., DOLSON, J., AND LEVOY, M. 2009. Gaussian KD-trees for fast high-dimensional filtering. *ACM Transactions on Graphics 28,* 3. Proceedings of the ACM SIGGRAPH conference.

BAE, S., PARIS, S., AND DURAND, F. 2006. Two-scale tone management for photographic look. *ACM Transactions on Graphics (Proc. SIGGRAPH) 25,* 3, 637 – 645.

BARASH, D. AND COMANICIU, D. 2004. A common framework for nonlinear diffusion, adaptive smoothing, bilateral filtering and mean shift. *Journal of Image and Video Computing 22*, 73–81.

BLACK, M. J., SAPIRO, G., MARIMONT, D. H., AND HEEGER, D. 1998. Robust anisotropic diffusion. *IEEE Transactions on Image Processing 7,* 3 (March), 421–432.

BUADES, A., COLL, B., AND MOREL, J.-M. 2005. Neighborhood filters and PDE's. Tech. Rep. 2005-04, CMLA.

BUADES, A., COLL, B., AND MOREL, J.-M. 2006. The staircasing effect in neighborhood filters and its solution. *IEEE Transactions on Image Processing 15,* 6, 1499–1505.

CHEN, J., PARIS, S., AND DURAND, F. 2007. Real-time edge-aware image processing with the bilateral grid. *ACM Transactions on Graphics (Proc. SIGGRAPH) 26,* 3.

DURAND, F. AND DORSEY, J. 2002. Fast bilateral filtering for the display of high-dynamic-range images. *ACM Transactions on Graphics (Proc. SIGGRAPH) 21,* 3.

ELAD, M. 2002. On the bilateral filter and ways to improve it. *IEEE Transactions On Image Processing 11,* 10.

FARBMAN, Z., FATTAL, R., LISCHINSKI, D., AND SZELISKI, R. 2008. Edge-preserving decompositions for multi-scale tone and detail manipulation. *ACM Transactions on Graphics (Proc. SIGGRAPH) 27,* 3.

FATTAL, R., LISCHINSKI, D., AND WERMAN, M. 2002. Gradient domain high dynamic range compression. *ACM Transactions on Graphics (Proc. SIGGRAPH) 21,* 3.

KASS, M. AND SOLOMON, J. 2010. Smoothed local histogram filters. *ACM Transactions on Graphics (Proc. SIGGRAPH) 29,* 3.

MRZEK, P., WEICKERT, J., AND BRUHN, A. 2006. *Geometric Properties from Incomplete Data*. Springer, Chapter On Robust Estimation and Smoothing with Spatial and Tonal Kernels.

PARIS, S. AND DURAND, F. 2009. A fast approximation of the bilateral filter using a signal processing approach. *International Journal of Computer Vision*.

PARIS, S., HASINOFF, S. W., AND KAUTZ, J. 2011. Local Laplacian filters: Edge-aware image processing with a Laplacian pyramid. *ACM Transactions on Graphics (Proc. SIGGRAPH) 30,* 4. to appear.

input image (HDR, gamma compressed)

model

our method

Sunkavalli et al. [2010]

Naive gradient transfer using a Poisson equation

Bae et al. [2006]

Fig. 11: *In contrast to other methods, our style transfer method is able to reveal a lot of small variations in the input image.*

PARIS, S., KORNPROBST, P., TUMBLIN, J., AND DURAND, F. 2009. Bilateral filtering: Theory and applications. *Foundations and Trends in Computer Graphics and Vision*.

PERONA, P. AND MALIK, J. 1990. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions Pattern Analysis Machine Intelligence 12,* 7 (July), 629–639.

SUBR, K., SOLER, C., AND DURAND, F. 2009. Edge-preserving multi-scale image decomposition based on local extrema. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia) 28,* 5.

SUNKAVALLI, K., JOHNSON, M. K., MATUSIK, W., AND PFISTER, H. 2010. Multi-scale image harmonization. *ACM Transactions on Graphics (Proc. SIGGRAPH) 29,* 3.

TOMASI, C. AND MANDUCHI, R. 1998. Bilateral filtering for gray and color images. In *Proceedings of the International Conference on Com-*

Fig. 12: *In the method of Bae and colleagues, the Poisson reconstruction is computationally expensive, and often breaks the style transfer while trying to fix the strong halos. In contrast, our direct method does not create halos and achieves acceptable style transfer even on these difficult examples.*

puter Vision. IEEE, 839–846.

VAN DE WEIJER, J. AND VAN DEN BOOMGAARD, R. 2002. On the equivalence of local-mode finding, robust estimation and mean-shift analysis as used in early vision tasks. In *Proceedings of the International Conference on Pattern Recognition*.