Vis Comput (2010) 26: 739–748
DOI 10.1007/s00371-010-0485-3

ORIGINAL ARTICLE

Contrast prescription for multiscale image editing

Dawid Pająk · Martin Čadík · Tunç Ozan Aydın · Makoto Okabe · Karol Myszkowski · Hans-Peter Seidel

Published online: 14 April 2010

Abstract Recently proposed edge-preserving multi-scale image decompositions enable artifact-free and visually appealing image editing. As the human eye is sensitive to contrast, per-band contrast manipulation is a natural way of image editing. However, contrast modification in one band usually affects contrasts in other bands, which is not intuitive for the user. In practice, the desired image appearance is achieved through an iterative editing process, which often requires fine tuning of contrast in one band several times. In this article we show an analysis of properties of multiscale contrast editing frameworks and we introduce the concept of contrast prescription, which enables the user to lock the contrast in selected areas and bands and make it immune to contrast manipulations in other bands.

D. Pająk (✉)
Computer Science Department, West Pomeranian University of Technology, Żołnierska 49, 71-210 Szczecin, Poland
e-mail: dpajak@mpi-inf.mpg.de

M. Čadík · T.O. Aydın · M. Okabe · K. Myszkowski · H.-P. Seidel
Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany

M. Čadík
e-mail: mcadik@mpi-inf.mpg.de

T.O. Aydın
e-mail: tunc@mpi-inf.mpg.de

M. Okabe
e-mail: mokabe@mpi-inf.mpg.de

K. Myszkowski
e-mail: karol@mpi-inf.mpg.de

H.-P. Seidel
e-mail: hpseidel@mpi-inf.mpg.de

Keywords Multiscale image editing · Contrast enhancement · Interactive image processing · HDR · Computational photography · Image decomposition

Fig. 1 An application of prescription idea in multiscale contrast manipulation. Enhancement of medium/high frequency contrast bands produces saturation of some image features existing in lower bands. Prescribing the contrast of unmodified bands prevents the saturation and at the same time allows to effectively increase the contrast according to user's request

1 Introduction

Contrast editing is a common post-processing step in digital photography, usually aimed towards improving the photograph's aesthetical appeal. Research on contrast editing has been focused on two main issues: developing versatile, yet computationally efficient frameworks that produce artifact-free results, and developing user interfaces that provide intuitive interaction with these underlying frameworks. Common to all state-of-the-art contrast editing methods is the involvement of a multiscale image decomposition, through which the image contrast can be edited in an arbitrary number of scales. This tendency is not surprising; since the human visual system (HVS) comprises mechanisms to perceive contrast in multiple spatial frequencies, editing fine and coarse image details separately feels only natural for the end user.

In a multiscale framework, a perfect separation between individual scales such that no two scales have any overlap, can theoretically be achieved by using frequency domain filters with sharp cutoffs. However this simple approach is never applied in practice since it results in heavy ringing artifacts, which can only be avoided by a smooth transition between filters of different scales. As a direct result, an important, but often ignored property of multiscale frameworks is the interaction between contrast at individual scales. Thus, enhancing the contrast of, for example, medium frequency details, indirectly affects the appearance of fine and coarse details due to the overlap between the filters of the neighboring scales in frequency domain. The central idea of this work is "contrast prescription", where the user selects a certain image region for which the contrast at each unmodified scale is locked ("prescribed"), and thus the image details with the desired spatial frequency can be edited independently.

The practical implication of contrast prescription is a more intuitive contrast editing experience. While a set of user controls (often sliders) that control the contrast amplitude at different scales gives the impression of being orthogonal to each other, in reality the changes at one scale propagates to others; an effect which we call "leaking" of contrast. These leaks can effectively be prevented by prescribing the contrast in the image region being edited, which frees the user from iteratively adjusting interface controls to confine the contrast change to the desired scale.

We show that contrast prescription can be implemented in multiple state-of-the-art contrast editing frameworks. Our GPU implementation combined with an intuitive user interface comprising brush and slider controls provides real-time feedback. In this work we focus on editing both low-dynamic range (LDR) and high-dynamic range (HDR) images using ordinary (LDR) display devices. In the rest of the paper we discuss related work on contrast editing and multiscale image decompositions (Sect. 2), give details on the related consequences of multiscale editing (Sect. 3) and how we address them (Sect. 4). Next, we discuss details on the extension of multiple frameworks to handle contrast prescription (Sect. 5), and finally present our results (Sect. 6).

2 Related work

Multiscale image decompositions such as the Laplacian pyramid [3, 11] have been successfully applied to many image processing tasks, including image editing. The practical advantage of considering multiple scales for image editing is the ability to modify the appearance of coarse and fine details separately [17]. On the downside, enhancing fine details disproportionally to coarser details leads in the extreme case to the well-known halo artifacts, resulting in unnatural images often considered not to be aesthetically pleasing.

Edge-preserving image decompositions, on the other hand, minimize halo artifacts by avoiding smoothing across strong edges. The edge-preserving behavior is accomplished through non-linear filters such as weighted least squares [14], anisotropic diffusion [2, 24], or the bilateral filter [27]. Motivated by the anisotropic diffusion, [28] proposed a hierarchical approach called LCIS for HDR tone mapping purposes. The bilateral filter has been widely used in HDR tone mapping as well [5], but also in image fusion [6, 9], example-based transfer of photographic look [1], among others. However, the extra complexity of the filters confine applications of Bilateral filtering to work offline. The performance issue has been addressed by introducing the "Bilateral grid" as a data structure built on Bilateral filter, which enables real-time, multiscale edge-aware image manipulations [4]. The edge-preserving behavior has been further improved by a weighted least squares (WLS) based framework [7], and later another framework based on edge avoiding wavelets [8] has been shown to achieve similar quality results much faster, due to the involvement of the linear time lifting scheme. The principle idea of preserving edges during decomposition has also been used in contrast processing of HDR images [21]. This method relies on performing image editing by first scaling perceptually linearized image gradients, and then reconstructing the image from the new gradient values. Recently, Subr et al. [25] proposed another edge-preserving image decomposition based on local extrema.

One consequence of multiscale image editing is the increased complexity of the editing process from the user's perspective. In fact, interfaces for intuitive image manipulation have been an active topic of research [15, 16, 18]. In our work we used an intuitive brush based interface, using which we were able to generate results in the paper in sessions lasting a few minutes. This was achieved also thanks to the real-time feedback of the underlying contrast processing framework.

3 Consequences of band modification in multiscale image decomposition

Contrast can vaguely be described as the difference between the intensity of an image location with the intensity of some neighborhood around that location, normalized again by the intensity of the same neighborhood. A mathematical formulation of this quantity is possible for simple luminance patterns (such as a foreground–background stimulus with luminance profile defined by a step function) where contrast can be defined as Weber's ratio:

$$W = (L_{\max} - L_{\min})/L_{\min}, \tag{1}$$

or the logarithmic ratio

$$G = \log(L_{\max}/L_{\min}), \tag{2}$$

where $L_{\max}$ and $L_{\min}$ are the intensities of the foreground and the background, respectively. For a sinusoidal luminance pattern, contrast can be computed using Michelson's formula $M = (L_{\max} - L_{\min})/(L_{\max} + L_{\min})$ with $L_{\max}$ and $L_{\min}$ being the sinusoid's peak points. Note that choosing among these "simple" contrast definitions is contextual, since they can trivially be converted to one another if required.

Natural images, however, are much more complex than mere step or sinusoidal intensity patterns, in that they contain various details at multiple scales. Consequently, the computation of the aforementioned "simple" contrast measures is not clear since $L_{\max}$ and $L_{\min}$ are not well-defined. Peli [23] defines contrast in complex images as the ratio of the bandpass image to the low-pass image at multiple scales

$$C_i = \frac{K_{\sigma(i)} * I - K_{\sigma(i+1)} * I}{K_{\sigma(i+1)} * I}, \tag{3}$$

where $*$ denotes the convolution operation between linear luminance and a low-pass Gaussian kernel $K_{\sigma(i)}$ at scale i, and $\sigma(i) = 2^i/\sqrt{2}$ denotes standard deviation, which accounts for frequency band cutoff.

In this paper, we use a multiscale contrast representation, where each contrast sub-band is calculated as a ratio between successive (i and $i+1$) Gaussian-like[1] smoothings of the image I:

$$G_i = \frac{\text{smoothingOperator}(I, i)}{\text{smoothingOperator}(I, i+1)}. \tag{4}$$

To simplify the computations, the decomposition is performed on the logarithm of luminance (roughly approximating the non-linear perception of luminance), for which the ratio can be replaced with a simple subtraction. Such a difference then gives the logarithmic ratio between an image location at some scale and the mean of its neighborhood, as shown in (2). This representation has also the advantage of being computationally more efficient then Peli's contrast, therefore most of the multiscale image editing frameworks [8, 19] follow this simple band decomposition scheme.

The *selection of smoothing operator* is usually the key factor in the overall performance and quality of a decomposition framework. Marr and Hildreth [22] define two main requirements that need to be met for the good smoothing filter. The first is that every multiscale decomposition requires the filter spectrum to be smooth and band-limited in the frequency domain. This allows to reduce the range of scales over which intensity changes take place. We can design a band-pass filter that would be perfectly localized in the frequency domain (sharp cutoff or brick-wall type of filter). However, processing an image with such a filter will induce well known ringing artifacts. Non-oscillating low-frequency parts of the image will yield in global oscillations in the output band representation. To prevent such artifacts, one needs to put the second requirement, a spatial localization constraint on the filter characteristics. This requirement, much more important from image editing point-of-view, can be interpreted in a way that every pixel in the filtered image should be computed from a weighted average of nearby pixels. The constraints above are contradicting in a sense that we are able to increase the spatial locality (reduce ringing) at the cost of frequency domain performance (reduced band-pass behavior).

A good example of such a trade-off in filter design is the Gaussian low-pass filter. It is non-negative and non-oscillatory, hence causes no ringing. The response in the frequency domain is a Gaussian function itself with the mean focused around the middle frequency of the band. This feature has an important consequence when applied in multiscale image processing. Such filter is inevitably causing an *energy leakage* between bands: the energy (modification) in one band leaks out to neighboring bands in the successive multiscale image manipulations. We show an illustrative band manipulation example exploiting Gaussian filter in Fig. 2 (left column), where single sub-band modification generates undesired energy leaking in neighboring sub-bands (Fig. 2, middle column). The energy leakage is prevented when ideal band-pass filter is used (identified by box function in frequency domain), however it results in ringing artifacts as described above. In Fig. 3 we compare the energy leakage of Gaussian-based image decomposition with this "ideal" band-pass decomposition.

The uniform smoothing behavior of the Gaussian filter, which leads to halo artifacts if sub-bands are independently modified, is addressed by so-called edge-preserving smoothing operators. They preserve sharp edges by excluding pixels across image discontinuities from consideration, which

[1] In practice, any type of low-pass (and also edge-preserving) filter can be used as a replacement of Gaussian filter.

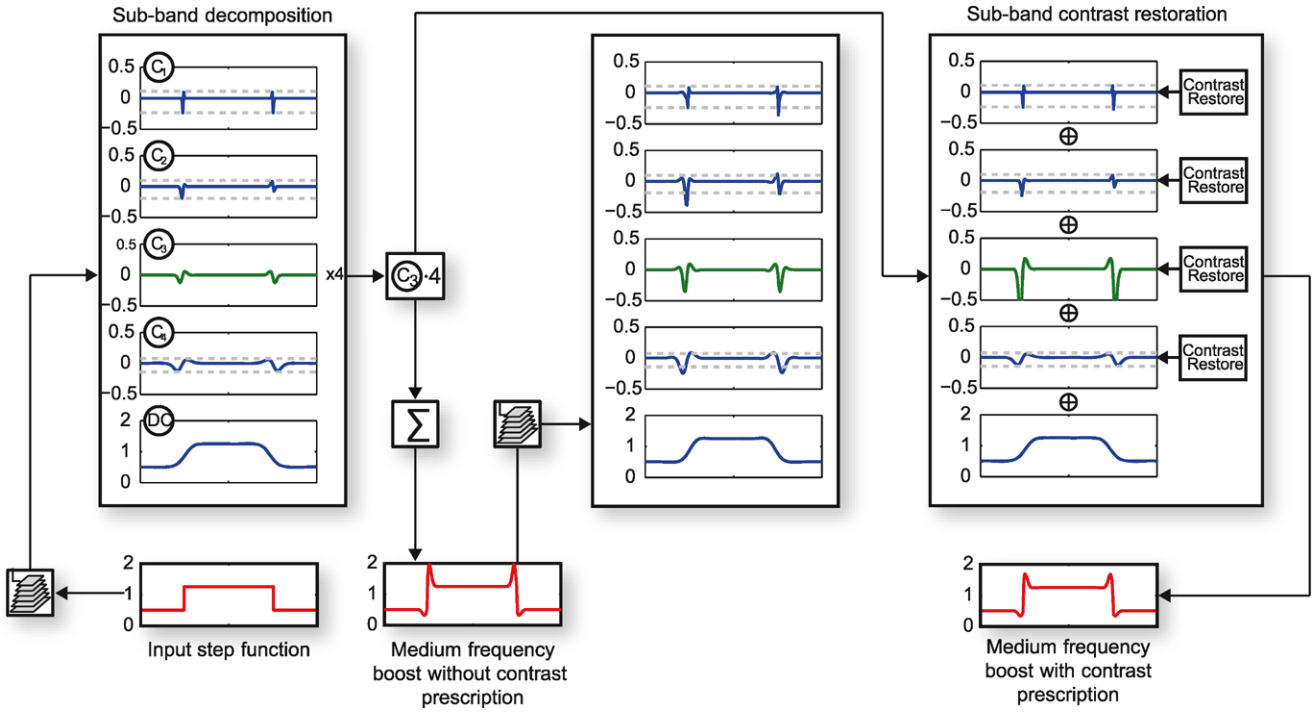

Fig. 2 Step function contrast enhancement. Input function is decomposed into 4 frequency bands and a DC component. Each band is visualized using Peli's definition for the physical contrast. We simulate detail enhancement by multiplying the C_3 band by 4. Decomposing the modified signal again shows that the contrast change of modified band is not proportional to the applied multiplier (*middle column*). Furthermore, due to Gaussian filter *energy leaking* property, all neighboring bands have been modified. Contrast prescription during the computation of modified signal (*right column*), *locally* limits the energy leaking and allows for better preservation of the contrast values in unaffected bands

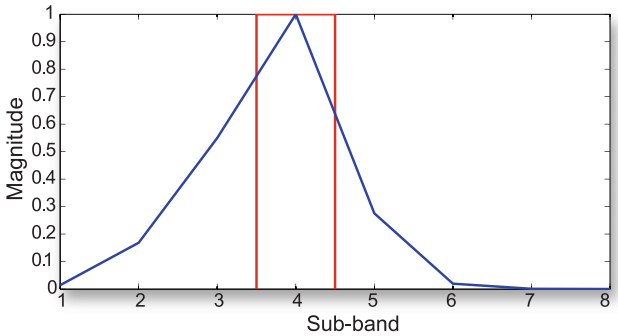

Fig. 3 Comparison of *energy leakage* of a Gaussian based decomposition with an ideal band-pass filter based decomposition. Input impulse signal is decomposed into 8 DoG (Difference-of-Gaussians) bands. While modifying the band we measure the signal magnitude change at each scale, which leads to piecewise linear approximation (*blue plot*). Due to non-ideal band-pass characteristics of the Gaussian filter, boosting the *middle band* results in uncontrolled amplification of features in neighboring scales (*blue plot*). The energy distribution of DoG decomposition resembles Gaussian function itself, but is not symmetric in logarithmic frequency scale. On the contrary, ideal band-pass filter satisfies the localization requirement in frequency domain (*red*). However, when applied to discrete image edges, creates undesired ringing effect in the spatial domain

avoids dividing the energy of the same edge across multiple sub-bands. However, due to imperfect localization in the spatial and frequency domain of the low-pass filter, these halo-free decompositions are still affected by the inter-band energy leakage during band manipulation.

We are not aware of any practical image decomposition scheme which is free of inefficiencies of Gaussian-like low-pass filtering mentioned here. To address this problem we introduce *contrast prescription* which restores unmodified bands physical contrast (i.e. Peli's contrast) during image reconstruction process (see Fig. 2, right column).

4 Contrast prescription

To overcome the inconvenient effect of energy leakage in multiscale based contrast editing applications, we introduce a simple and efficient contrast prescription idea. Our proof-of-concept implementation allows to directly manipulate spatial selection of entire range of image scales. We define prescription-enabled editing as a manipulation in which all unmodified sub-bands are prescribed to keep their contrast values constant. As there are no assumptions made on the type of chosen multiscale decomposition method, the contrasts in prescribed sub-bands are restored while the pyramid-like decomposition is integrated back to the output

image. Adopting contrast representation from (4), the integration is performed by simple addition of all sub-bands:

$$\log(I_{\text{out}}) = \sum_{i=1}^{N} \text{mult}_i \cdot G_i, \tag{5}$$

where mult_i denotes per pixel contrast multiplier for band i and $\cdot$ is an element-wise multiplication operator. Note that multipliers are applied to the logarithmic contrast representation, which is equal to computing a power function on the linear contrast values. In practice, this prevents too strong darkening of the image which otherwise would happen in linear space. However such an operation, in case of detail enhancement, tends to oversaturate the already well-visible details. Our prescription algorithm can counteract this scenario by selectively modulating the contrast values of oversaturated pixels. The overall contrast restoration algorithm is shown as pseudo-code in Fig. 4.

Although we store sub-band contrast in a logarithmic ratio representation, in order to compute contrast changes during successive band manipulations, we utilize Peli's physical contrast measure, as it is a metric that can be reliably applied to complex images and by definition takes into account inter-band dependencies. We employ local sub-band physical contrast ratio to correct the contrast values affected by the cross-band energy leaking. The contrast correction algorithm starts from the lowest frequency band. According to (3), if the band-pass image is constant, the only way to change contrast is to modify the low-pass image of the same band. As we perform interactive multiscale contrast manipulation, the aforementioned situation takes place quite often. In order to correct for this, we simultaneously compute two low-pass images I_0 and I_1 (approximated background luminance), which let us estimate for each sub-band how the contrast has changed. Initially both images are the same, so no correction is applied. However, as we move further, adding up a new sub-band components (Line 9–10 of Fig. 4), the difference between background luminance values becomes more apparent, and directly affects the Peli contrast for prescribed bands. We modify the *mult* set to reflect the background luminance change and locally correct for the suppression or enhancement of prescribed sub-band contrast.

Given a *linear* low-pass image I_1 and its prescribed counterpart I_0, we perform pixel-wise scaling of mult (line 7) by $(\frac{I_0}{I_1})^{\beta}$ ratio, which for $\beta = 1$ corresponds to restoring Peli's contrast for a certain pixel. The β scaling parameter provides a non-linear control over the performance of contrast restoration. As the sub-band contrast is stored in a low-pass logarithmic format, before changing the contrast value we need to convert it to an intermediate notation for which the linear band-pass signal is expressed in the nominator. For this purpose, we use the Weber fraction computed as $W = \Delta L/L = 10^{|G|} - 1$. After modifying the contrast we apply simple transformation to get back to low-pass logarithmic contrast: $G = \text{sign}(G) \cdot \log(W + 1)$. Corrected mult multipliers are stored in a separate location, so they can be used in construction of reference contrast pyramid G in user's future edits. After applying the contrast restoration formula for all sub-bands we output the image to the display.

The bottom-up approach is motivated by the fact that the most of the energy leaking is due to the large signal amplitudes of low-frequency bands. This is supported by the findings in natural image statistics [10]: most of the image energy is focused around low-frequency bands, which is known as the power law for the amplitudes of frequencies.

1: **ContrastRestore**(MultiScaleMultipliers mult, MultiScaleDecomposition G)
2: $N \leftarrow \text{height}(G)$ // sub-band count
3: $I_0 \leftarrow G_N$ // log(DC)
4: $I_1 \leftarrow G_N$ // log(DC)
5: **for** $i \leftarrow N-1$ **downto** 1 **do**
6: **for all** pixels **do**
7: $R \leftarrow \text{mult}_i \cdot 10^{(I_0 - I_1)\cdot\beta}$ // corrected multiplier
8: $W \leftarrow 10^{|G_i|} - 1$
9: $I_0 \leftarrow I_0 + G_i$
10: $I_1 \leftarrow I_1 + \text{sign}(G_i) \cdot \log(W \cdot R + 1)$
11: **end for**
12: **end for**
13: $I_{\text{out}} \leftarrow I_1$

Fig. 4 Contrast restoration algorithm. We simultaneously integrate two pyramids using previous and current band multipliers. The integration incorporates a correction fraction for I_1 (output) that locally restores contrast in prescribed bands. The I_0 image serves as a prescribed sub-band adaptation luminance reference which is required for computing the correction factor

5 Extension of edge-preserving decompositions

In this section we discuss extending recent edge-stopping multiscale decomposition frameworks with contrast prescription. Furthermore, we introduce supplementary extension that allows the user for counter-shading (halo editing).

5.1 Weighted least squares decomposition

The contrast prescription algorithm can be implemented in WLS optimization framework in a straightforward manner. The basic idea behind the WLS based decomposition is to keep the complete frequency domain representation of each edge at only one scale. The multiscale image decomposition is achieved by iterative application of an edge-stopping image smoothing operator, which is tuned in subsequent iterations to preserve edges of successively larger contrast. In order to convert such a representation to contrast sub-bands,

we employ (4). The smoothing operator is designed as a Poisson-class linear optimization which minimizes the energy function that penalizes the image gradients (smoothing effect) in the whole image except near strong edges (edge-stopping effect). Due to the existence of high frequency edges in the band-pass image, the WLS filter requires the bands to be stored as full-resolution images.

The weighting function, responsible for edge-stopping behavior is expressed as:

$$w_n(m) = \frac{1}{|L_n - L_m|^{\alpha} + \epsilon}, \tag{6}$$

where w_n denotes an edge weight between pixel value L_n and its neighbor L_m (α is a model parameter and ϵ prevents division by zero). Contrast restoration mechanism (see Fig. 4) is applied directly on the WLS multiscale image decomposition and the resulting image is obtained by adding up the sub-bands.

5.2 Second generation wavelet decomposition

Recent work [8] showed the application of second generation wavelets to edge-preserving image decomposition. Here, the image is decomposed using edge avoiding wavelets (EAW)—second generation wavelets constructed with a weighting function similar to (6). The computation of second generation wavelet decomposition is performed using the lifting scheme [26].[2]

Our wavelet implementation is based on a Weighted Red-Black decomposition (WRB) [29]. After transforming the image into wavelet representation we cannot directly use the wavelet scaling function coefficients to apply our contrast prescription algorithm. In order to recover DoG-like decomposition of the image, we compute N inverse wavelet transformations. Each inverse transformation sets all scaling coefficients to zero, except the ones which describe features at scale i. Such an algorithm performs an edge-aware interpolation (upsampling) of selected sub-band components. The output image is a full-resolution sub-band which roughly corresponds to the results obtained by the WLS framework mentioned above. As we show in Sect. 6 our GPU implementation of wavelet decomposition is very fast; the entire process takes less than 5 ms on mainstream hardware.

5.3 Perceptual contrast processing framework

Mantiuk et al. [21] presents a framework in which the inter-band dependencies are tightly integrated in the inhomogeneous Laplace equation and the prescription algorithm cannot be applied in the form described earlier. In this framework, especially suitable for processing HDR images, the final image is a result of least square optimization (computed using a Poisson solver) and is not reproduced by simple addition of sub-bands. In order to make our approach applicable we implement contrast restoration as a post-processing step. We use two separate decompositions to track the changes before and after manipulations, constructed by the EAW algorithm described earlier due to its efficiency. Contrast prescription is then realized using sub-band contrasts obtained from these external decompositions.

5.4 Interactive halo editing

The use of counter-shading to enhance the perceived contrast has been known by painters for ages [20]. More recently, [13] proposed an automatic technique for improving contrast perception in digital images by modulating brightness at the edges. In our technique, we allow the user to control the halo effect (counter-shading) *manually*. This is implemented inside the edge-preserving decomposition by means of a minor modification in the weighting function (6) used by both decomposition frameworks. By modulating α coefficient, which is responsible for edge-stopping behavior of decomposition scheme, we can suppress or enhance halo effect near the edges. When α is close to 0, the weights are becoming more spatially uniform, thus the smoothing operator resembles regular Gaussian-like filter. This results in a decomposition that, in case of a local manipulation, is affected by the halo effect. For each sub-band, we define the α pixel-wise. To modify these coefficients we use the same approach as for updating contrast multipliers (see Sect. 6).

6 Results

In this section we present results of a comparison of prescription-enabled editing with regular one on a number of images. Our proof-of-concept software[3] was tested on a mainstream PC equipped with Intel Core2 Duo 3.0 Ghz CPU and NVidia GTX260 GPU. We compared the performance of decompositions presented here. In most cases the decomposition can be done off-line, as a preprocessing step before actual editing session. However, features like halo editing require recreating the sub-bands every time we modify edge weights. Consequently, we chose wavelet decomposition as our benchmark implementation since it is significantly faster than other schemes and the results we obtained are comparable. For a 1 MPixel image, the forward wavelet transform coupled with generation of 8 full-resolution contrast bands takes less than 5 ms. Hence, the framework runs at interactive speeds even

[2] We refer the reader to Jansen and Oonincx [12] for a detailed discussion on second generation wavelets.

[3] The implementation comprises of a platform independent Java UI and native GLSL image processing library.

```
1: UpdateMultipliers(MultiScaleMultipliers mult, Image mask)
2: N ← height(mult)  // sub-band count
3: for i ← 1 to N − 1 do
4:    B = GetBandMultiplier(i)
5:    for all pixels do
6:       mult_i ← max(0, mult_i + mask · B)
7:    end for
8:    mask = mask ↓ 2
9: end for
```

Fig. 5 Contrast multipliers update. The GetBandMultiplier(i) function returns [0, 1] normalized value which indicates mask scaling factor for band i. It can be either user-defined by setting up scale range sliders or computed in fully automatic manner

for large images. In case of Poisson solver, on the other hand, each iteration takes about 3 ms, depending on the amount of modification applied. Note that edits are performed iteratively on a small parts of the image. Therefore, the solver is initialized with a good quality solution, which only needs to be corrected in some selected regions. On average our conjugate gradient based solver requires about 10 iterations to converge. The test software binaries are available for Windows/Linux and can be downloaded from (http://mpi-inf.mpg.de/~dpajak/prescription).

Our implementation allows the user to intuitively manipulate the contrast multipliers. Brushing over selected areas creates a smooth amplitude mask with Gaussian-like falloff, which is then used to update per band, per pixel multipliers (see Fig. 5). We decided to implement brush based interface as it is still considered to be the most common tool used for manual image retouching. However, this interface can be easily extended by introducing more automated, diffusion-based segmentation as in [19]. For a novice user, manipulating bands with scale range sliders can be challenging. Therefore, we include a brushing mode where band multipliers (see Fig. 5) are computed automatically by measuring image energy[4] for current selection. This approach will always try to selectively boost the bands with smallest energy value.

6.1 Contrast prescription

We demonstrate our approach by performing a set of exemplary, yet typical, contrast editing sessions. First, we show a simple scenario, where only one band range is modified and then we stage more complex manipulation to show how bands interact with each other.

In Figs. 1 and 6, we perform single modification of fine image details using WRB Wavelet decomposition scheme

[4]Modulo of a gradient for gradient domain frameworks and absolute amplitude of band-pass contrast for multi-scale decompositions.

Fig. 6 Example of contrast enhancement based on WRB wavelet decomposition for $\alpha = 0.8$ (6) and $\beta = 1.0$ (Fig. 4). Boosting fine/medium scales in the ceiling area causes saturation of some image features and results in unnatural appearance. Despite the extreme boost, contrast restoration allows to regain natural look of modified area and get the requested detail enhancement

Fig. 7 Comparison of regular and prescription-enabled image editing in WLS framework ($\alpha = 1.2$). Amplification of coarse features attenuates contrasts in higher bands. Contrast restoration algorithm successfully recovers fine details and boosts coarse image features at the same time

with parameters $\alpha = 0.8$ and $\beta = 1.0$. In both figures, we used the same band range multiplier for prescribed and non-prescribed operation. Due to the energy leakage, features existing across multiple scales are over saturated and by com-

paring against the source image we see that their contrast enhancement is spatially inconsistent. Prescription counteracts these situations and allows for more uniform and controllable contrast modification. Note that it is cumbersome to achieve such a result with a series of separate edits since the restoration is spatially local and highly non-linear.

Figure 7 shows an opposite scenario, where user scales low-frequency bands. The modification causes the loss of visibility of trees and plant pot details. Prescription of contrasts in upper bands allows to achieve both goals, increasing the global contrast and maintaining the detail visibility.

In Fig. 8, the image is initially modified by enhancing fine details around the plant area. Next, we boost low-frequency bands, which reduces the visibility of previous edits. Also, high signal amplitudes of low-frequency content exposes the energy leakage issue, resulting in an over-saturated image and decreased perception of unmodified contrasts bands. Contrast prescription visibly restores the details and reduces the saturation while still allowing for large enhancement of low-frequency contrasts.

Finally, we illustrate the complete editing session result, Fig. 10, where we manipulate contrasts in order to transfer the style of a professional photographer to a plain picture of the same location. Despite the obvious difference in source material we managed to properly reflect the style using only local, prescription-enabled contrast modifications.

Fig. 8 Contrast editing session in WLS decomposition framework ($\alpha = 1.2$). Contrast prescription prevents the loss of details in unmodified sub-bands, as a result, previously modified fine details are preserved

6.2 Interactive halo editing

As described in Sect. 5.4, we enhanced each of the implemented decomposition frameworks to allow interactive halo manipulation. Figure 9 illustrates simple, low-frequency halo suppression case. Although the local modification of edge-stopping filter behavior usually requires repeating the decomposition, we show that an efficient hardware implementation can deliver an interactive solution even in case of Poisson solver based frameworks.

7 Conclusions and future work

In this work, we analyzed the properties of multiscale image representations used in interactive image editing applications. Direct consequence of Gaussian-like filters commonly used to construct such representations (including edge-aware decompositions), is the effect that we call "energy leaking". When the user modifies one sub-band, part of the "energy" of this modification in effect leaks out to the other sub-bands in successive manipulations. This can affect the perception of already edited parts of the image and leads to non-intuitive, iterative way of image editing. Moreover, change in one band can result in oversaturation of another band (also possibly previously edited), e.g. when user boosts overall contrast, the tiny details might be lost. To overcome those limitations inherently imposed by all existing multiscale image editing frameworks, we propose the concept of contrast prescription, where once edited parts of the edited image retain their prescribed values. The aim is to restore the visibility of each sub-band contrast and limit the effects of cross-band energy leakage. Consequently, reduced number of decomposition related artifacts, allows for more intuitive and controllable multiscale contrast manipulation. To illustrate the concept, we show simple, but efficient and interactive extension of three state-of-the-art multiscale frameworks.

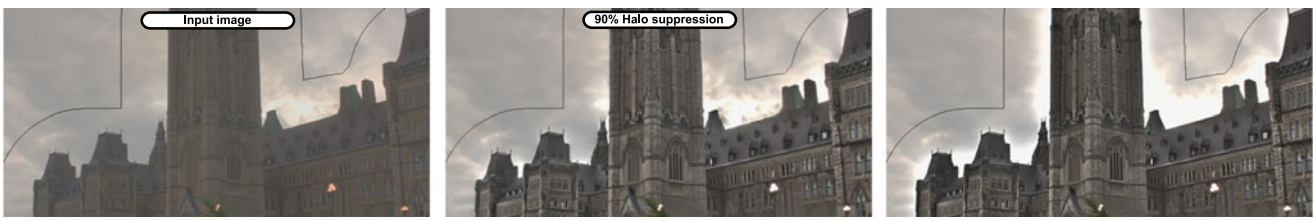

Fig. 9 Halo editing example with a Poisson solver based decomposition framework [21]. The original castle scene is modified by elevating high/medium frequency bands. To artificially modulate contrast along the edges, e.g. for better decomposition of foreground from background, one might locally allow the halo effect to be visible

Fig. 10 Practical demonstration of prescription-enabled contrast editing (*center*). We modified the source image (*left*) to reflect the style and feeling of pictures taken by Ansel Adams (*right*). As the artist style was based mostly on manual dodging and burning we had to perform extreme band manipulations that applied without prescription would result in heavy artifacts. The entire session was about 8 minutes long and included only contrast manipulations on WLS based decomposition

So far, we assumed editing using an ordinary (LDR) display device—in this case the illumination stayed almost constant and the effect of image editing on the user's visual adaptation was very subtle. However, when editing on an *HDR display* device, the change of user's adaptation due to the display luminance is not negligible any more and it can significantly bias user's perception. Modeling of apparent contrast is required to compensate for this effect which suggests a possible extension to this work.

References

1. Bae, S., Paris, S., Durand, F.: Two-scale tone management for photographic look. In: Proc. SIGGRAPH 2006, pp. 637–645 (2006)
2. Black, M., Sapiro, G., Marimont, D., Heeger, D.: Robust anisotropic diffusion. IEEE Trans. Image Process. **7**(3), 421–432 (1998)
3. Burt, P.J., Adelson, E.H.: The Laplacian pyramid as a compact image code. IEEE Trans. Commun. **31**, 532–540 (1983)
4. Chen, J., Paris, S., Durand, F.: Real-time edge-aware image processing with the bilateral grid. In: Proc. SIGGRAPH 2007, p. 103 (2007)
5. Durand, F., Dorsey, J.: Fast bilateral filtering for the display of high-dynamic-range images. In: SIGGRAPH '02: Proceedings of the 29th Annual Conference on Computer Graphics and Interactive Techniques, pp. 257–266. ACM, New York (2002)
6. Eisemann, E., Durand, F.: Flash photography enhancement via intrinsic relighting. ACM Trans. Graph. **23**(3), 673–678 (2004)
7. Farbman, Z., Fattal, R., Lischinski, D., Szeliski, R.: Edge-preserving decompositions for multi-scale tone and detail manipulation. ACM Trans. Graph. **27**(3) (2008)
8. Fattal, R.: Edge-avoiding wavelets and their applications. ACM Trans. Graph. **28**(3), 1–10 (2009)
9. Fattal, R., Agrawala, M., Rusinkiewicz, S.: Multiscale shape and detail enhancement from multi-light image collections. ACM Trans. Graph. (Proc. SIGGRAPH) **26**(3) (2007)
10. Gibson, J.D., Bovik, A. (eds.): Handbook of Image and Video Processing. Academic Press, Orlando (2000)
11. Gonzalez, R.C., Woods, R.E.: Digital Image Processing, 2nd edn. Prentice-Hall, New Jersey (2002)
12. Jansen, M.H., Oonincx, P.J.: Second Generation Wavelets and Applications. Springer, Berlin (2005)
13. Krawczyk, G., Myszkowski, K., Seidel, H.P.: Contrast restoration by adaptive countershading. Comput. Graph. Forum **26**(3), 581–590 (2007)
14. Lagendijk, R., Biemond, J., Boekee, D.: Regularized iterative image restoration with ringing reduction. IEEE Trans. Acoust. Speech Signal Process. **36**(12), 1874–1888 (1988)
15. Levin, A., Lischinski, D., Weiss, Y.: Colorization using optimization. SIGGRAPH **23**(3), 689–694 (2004)
16. Li, Y., Sun, J., Tang, C.K., Shum, H.Y.: Lazy snapping. SIGGRAPH **23**(3), 303–308 (2004)
17. Li, Y., Sharan, L., Adelson, E.H.: Compressing and companding high dynamic range images with subband architectures. ACM Trans. Graph. (Proc. SIGGRAPH) **24**(3), 836–844 (2005)
18. Li, Y., Adelson, E., Agarwala, A.: Scribbleboost: adding classification to edge-aware interpolation of local image and video adjustments. Comput. Graph. Forum **27**(4), 1255–1264 (2008)
19. Lischinski, D., Farbman, Z., Uyttendaele, M., Szeliski, R.: Interactive local adjustment of tonal values. ACM Trans. Graph. **25**(3), 646–653 (2006)
20. Livingstone, M.: Vision and Art: The Biology of Seeing. Harry N. Abrams (2002)
21. Mantiuk, R., Myszkowski, K., Seidel, H.P.: A perceptual framework for contrast processing of high dynamic range images. ACM Trans. Appl. Percept. **3**(3), 286–308 (2006)
22. Marr, D., Hildreth, E.: Theory of edge detection. Proc. R. Soc. Lond., Ser. B, Biol. Sci. **207**(1167), 187–217 (1980)
23. Peli, E.: Contrast in complex images. J. Opt. Soc. Am. A **7**(10), 2032–2040 (1990)
24. Perona, P., Malik, J.: Scale-space and edge detection using anisotropic diffusion. IEEE Trans. Pattern Anal. Mach. Intell. **12**(7), 629–639 (1990)
25. Subr, K., Soler, C., Durand, F.: Edge-preserving multiscale image decomposition based on local extrema. In: ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia 2009), Annual Conference Series. ACM, New York (2009)
26. Sweldens, W.: The lifting scheme: a construction of second generation wavelets. SIAM J. Math. Anal. **29**(2), 511–546 (1997)
27. Tomasi, C., Manduchi, R.: Bilateral filtering for gray and color images. In: Proceedings of the Sixth International Conference on Computer Vision (ICCV '98), p. 839. IEEE Computer Society, Los Alamitos (1998)
28. Tumblin, J., Turk, G.: LCIS: A boundary hierarchy for detail-preserving contrast reduction. In: Proc. SIGGRAPH '99, pp. 83–90 (1999)
29. Uytterhoeven, G., Roose, D., Bultheel, A.: Wavelet transforms using the lifting scheme (1997)

Dawid Pająk is a PhD student at West Pomeranian University of Technology. He received MSc Eng degree in computer science from Technical University of Szczecin in 2007. His research interests include image processing, rendering, high performance computing and GPGPU applications. Before moving into research he was a professional game developer.

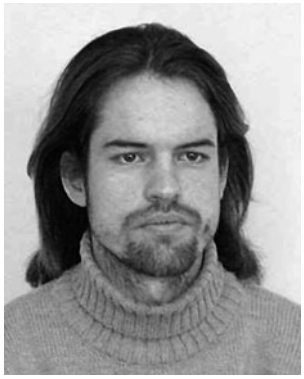

Martin Čadík is doing research in the fields of image and video quality assessment, high dynamic range imaging and image processing. He received MSc and PhD degrees in computer science from Czech Technical University in Prague in 2002 and 2008, respectively. He is currently a post-doc research fellow at Max-Planck-Institut für Informatik.

Tunç Ozan Aydın is a PhD student at Max-Planck-Institut für Informatik, and holds a BS degree in Civil Engineering from Istanbul Technical University and an MSc degree in Computer Science from Georgia Institute of Technology. His research interests are human visual system modeling, image/video quality assessment and HDR imaging.

Makoto Okabe is an assistant professor at department of computer science, the University of Electro-Communications. He received PhD from department of information science and technology, the University of Tokyo in March 2008. He worked as a postdoc in computer graphics group of Max-Planck-Institut für Informatik from 2008 to 2010. His research interest is user interface for computer graphics and current focus is on image and animation synthesis based on image and video database.

Karol Myszkowski is a senior researcher in the Computer Graphics Group of the Max-Planck-Institut für Informatik. From 1993 to 2000, he served as an Associate Professor at the University of Aizu, Japan. During the years 1985–1992, he worked as a Research Associate and then Assistant Professor at Szczecin University of Technology. He received his MSc degree in electrical engineering from Szczecin University of Technology, and PhD and habilitation degrees in computer science from Warsaw University of Technology in 1983, 1991 and 2001, respectively. His research interests include perception issues in graphics, high-dynamic range imaging, global illumination, rendering, and animation.

Hans-Peter Seidel is the scientific director and chair of the computer graphics group at the Max-Planck-Institut (MPI) Informatik and a professor of computer science at Saarland University. He has published and lectured widely. He has received grants from a wide range of organizations, including the German National Science Foundation (DFG), the German Federal Government (BMBF), the European Community (EU), NATO, and the German-Israel Foundation (GIF). In 2003 Seidel was awarded the 'Leibniz Preis', the most prestigious German research award, from the German Research Foundation (DFG). Seidel is the first computer graphics researcher to receive this award.