MIT Open Access Articles

*Edge-preserving Multiscale Image
Decomposition based on Local Extrema*

Citation: Kartic Subr, Cyril Soler, and Frédo Durand. 2009. Edge-preserving multiscale image decomposition based on local extrema. ACM SIGGRAPH Asia 2009 papers (SIGGRAPH Asia '09). ACM, New York, NY, USA, , Article 147 , 9 pages.

As Published: http://dx.doi.org/10.1145/1618452.1618493

Publisher: Association for Computing Machinery (ACM)

Persistent URL: http://hdl.handle.net/1721.1/72593

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Edge-preserving Multiscale Image Decomposition based on Local Extrema

Kartic Subr[*]
INRIA / Grenoble University

Cyril Soler[†]
INRIA / Grenoble University

Frédo Durand [‡]
MIT CSAIL

(a) Input	(b) Fine features boosted	(c) Coarse features boosted	(d) Scanline plots
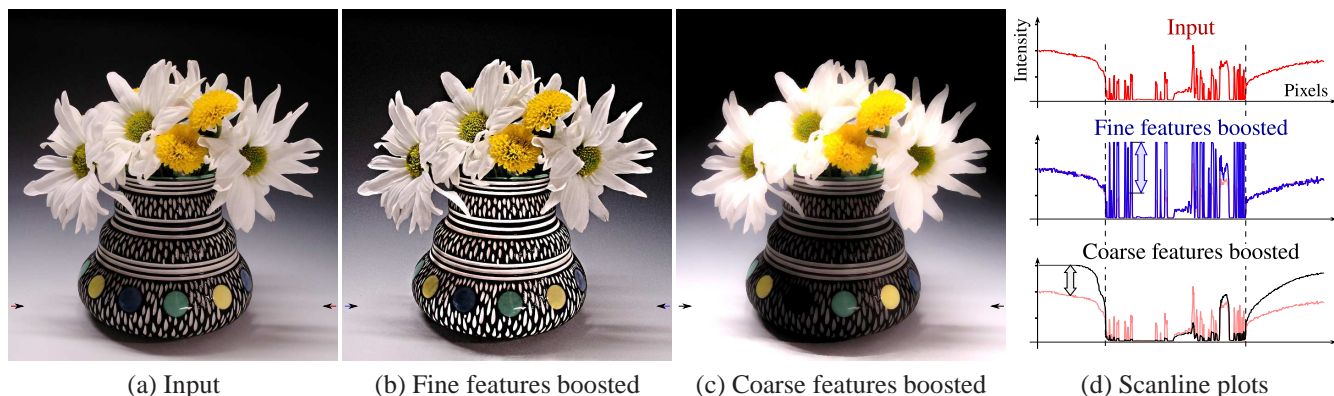
Figure 1: *Our multiscale decomposition of image (a) allows detail to be extracted based on spatial scale rather than contrast and preserves edges. (b) Boosting fine scale features increases the contrast of the pattern on the vase. (c) Boosting coarse scale contrast and suppressing fine features reduces the contrast of the pattern, while increasing the contrast of the vase with its background. (d) Scanline plots (rows indicated using arrows in (a), (b) and (c)), illustrating the effect of the two equalizations (b) and (c). The dashed lines in the plots show two examples of edges that have been preserved.*

Abstract

We propose a new model for detail that inherently captures *oscillations*, a key property that distinguishes textures from individual edges. Inspired by techniques in empirical data analysis and morphological image analysis, we use the local extrema of the input image to extract information about oscillations: We define detail as oscillations between local minima and maxima. Building on the key observation that the spatial scale of oscillations are characterized by the density of local extrema, we develop an algorithm for decomposing images into multiple scales of superposed oscillations.

Current edge-preserving image decompositions assume image detail to be low contrast variation. Consequently they apply filters that extract features with increasing contrast as successive layers of detail. As a result, they are unable to distinguish between high-contrast, fine-scale features and edges of similar contrast that are to be preserved. We compare our results with existing edge-preserving image decomposition algorithms and demonstrate exciting applications that are made possible by our new notion of detail.

Keywords: image decomposition, computational photography

1 Introduction

A variety of applications in computational photography require a decomposition of an image into different scales. Traditional approaches that use linear bases have evolved to accommodate the need for respecting strong edges. Recent definitions of scales are usually based on spatial scale definitions combined with a notion on the range to differentiate strong edges [Tomasi and Manduchi 1998; Durand and Dorsey 2002; Farbman et al. 2008; Lischinski et al. 2006; Choudhury and Tumblin 2005]. Current approaches share a common notion of an edge– large gradients, or large value differences, where the definition of large might depend on the application. However, this notion of an edge makes it challenging to capture fine details or textures that have fine spatial scale but high contrast. For example, in Figure 1(d), some edges to be preserved are lower contrast than oscillations to be smoothed. Extracting the white dots on the vase as detail requires aggressive smoothing of gradients, which would also blur single edges that are to be preserved (see Fig. 2). This distinction between edges and oscillations raises challenges in defining fully multiscale decompositions because the interplay between spatial and edge consideration leads to unexpected results, as shown by Farbman et al. [2008]

We propose a novel non-linear image decomposition that effectively extracts fine-scale features, regardless of their contrast, as detail and yet preserves softer salient edges in the base layer. In contrast to previous approaches that rely on magnitudes of pixel differences at their heart, our approach captures local image oscillations by considering local image extrema. A fine-scale texture is characterized by rapid oscillations (see Fig. 1) between minima and maxima. Furthermore, the oscillation between extrema provide critical information that permit the distinction of individual edges from oscillations. We obtain a multiscale decomposition by recursively smoothing the image while also progressively coarsening the scale at which extrema are detected.

1.1 Related work

Several image decomposition techniques have been proposed. Strategies that use linear filters [Burt and Adelson 1983; Rahman and Woodell 1997; Pattanaik et al. 1998] produce halo artifacts at edges and have been succeeded by non-linear filters that preserve strong edges– a popular choice being the bilateral filter [Tomasi and Manduchi 1998; Durand and Dorsey 2002; Choudhury and Tumblin 2005]. Bae et al. [2006] used the bilateral filter to separate images into low- and high-contrast features and manipulated the layers independently to enhance photographic look. Fattal et

[*]e-mail: Kartic.Subr@inrialpes.fr
[†]e-mail: Cyril.Soler@inrialpes.fr
[‡]e-mail: fredo@mit.edu

al. [2007] presented a technique to enhance shape and surface details of objects using bilaterally filtered representations of a set of differently lit images. Our goal is to extract from a single image, at each scale, the finest spatial oscillations as detail without assuming them to be low-contrast oscillations.

Two approaches have been proposed for multiscale decompositions using the bilateral filter. One strategy is to progressively increase the width of the range and spatial Gaussian through the coarsening process. Chen et al. [2007] used this technique to construct a *bilateral pyramid* for progressive video abstraction. Another strategy [Fattal et al. 2007] recursively applies the bilateral filter to the smoothed versions of the input image. This strategy decreases the width of the range-Gaussian during successive iterations so that edges from preceding smoothing operations are not blurred during the coarsening.

In recent work, Farbman et al. [2008] pointed out that, while the bilateral filter is effective at smoothing out low amplitude noise at a fine scale, multiscale decompositions using the bilateral filter suffer from a variety of problems. Progressive widening of the range and spatial Gaussians through the coarsening process was shown to produce halo artifacts at strong edges. To overcome some problems of using the bilateral filter in a multiscale decomposition, Farbman et al. [2008] proposed a filter that smoothes an input image I by computing an image that is as close to I as possible while being smooth everywhere except at regions where the gradient of I is large. They used a weighted least squares filter, originally used to control ringing during deblurring of noisy images [Lagendijk et al. 1988]. The nature of this optimization makes it impossible to preserve salient edges with lower contrast than the texture that is to be smoothed.

In summary, smoothing filters currently used in image decomposition algorithms assume detail is low-contrast. As a result, local variation at different contrast levels are extracted as successive layers of detail. Such layers of detail do not necessarily represent fine-scale spatial variation.

A notable exception, for 1D data, is empirical mode decomposition [Huang 1998]— a powerful data analysis tool originally proposed to decompose nonlinear, nonstationary signals into their intrinsic modes of oscillations. The decomposition is achieved by iterative removal of the finest intrinsic oscillations as indicated by local extrema. This technique is popularly used on 1D data that do not contain sharp discontinuities. A few attempts at extending the technique to image decomposition [Nunes et al. 2003; Liu and Peng 2005; Damerval et al. 2005] have uncovered a number of difficulties. One formidable challenge that has not been addressed is the need to respect sharp edges. Another drawback of empirical mode decomposition is its poor handling of signals where oscillations at different scales occur as bursts, in parts of the domain (the problem of intermittency [Li et al. 2005]).

1.2 Contributions

We introduce novel definitions, based on local extrema, for edges and detail that permit the distinction between highly contrasted texture and single edges. Using these definitions we develop an edge-preserving smoothing algorithm that allows fine scale detail to be extracted regardless of contrast. We perform an edge-preserving multiscale decomposition by recursively applying the smoothing algorithm on the base layer. The decomposition corresponds to features at different spatial scales with salient edges being preserved. We compare our approach with existing decompositions and demonstrate its effectiveness using applications. Figure 4 places our novel algorithm in the context of existing approaches.
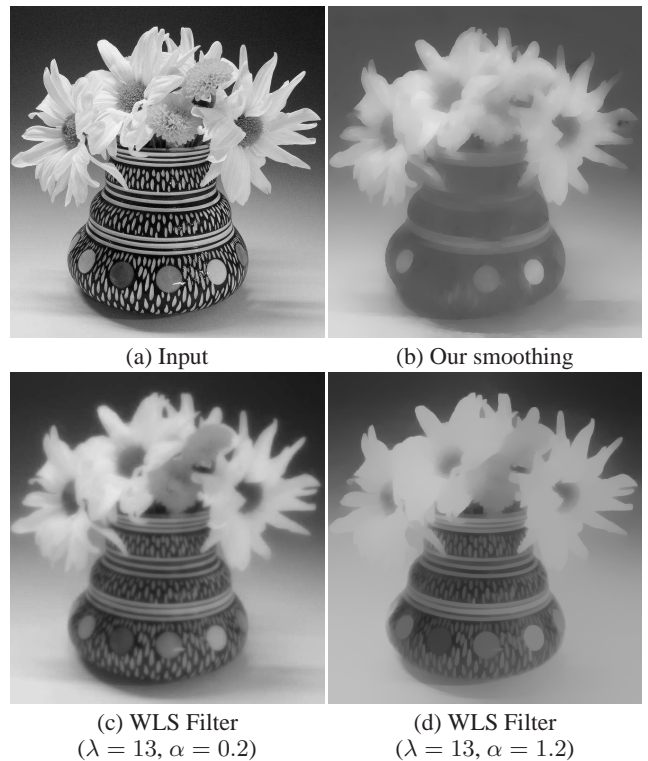

(a) Input   (b) Our smoothing

(c) WLS Filter   (d) WLS Filter
($\lambda = 13, \alpha = 0.2$)   ($\lambda = 13, \alpha = 1.2$)

Figure 3: *The ubiquitous notion of edges as pixels with large gradients does not allow disambiguation between fine-scale features and edges that are to be preserved, as shown by this example.* (a) *The contrast of the pattern on the flower vase is greater than across the edges of the soft shadows and petal boundaries.* (b) *Using our smoothing algorithm, the pattern is extracted as detail because of its fine scale, while coarser soft shadow- and petal-boundaries are preserved.* (c) *The weighted least square (WLS) filter does not smooth the pattern if fidelity to strong gradients is retained.* (d) *On the other hand, the WLS filter necessarily blurs softer edges even though they are coarse-scale features while smoothing the pattern on the vase.*

2 Extrema-based multiscale decomposition

We present a novel smoothing algorithm which effectively smoothes highly contrasted oscillations while preserving salient edges. By applying this algorithm recursively on the smoothed image, we compute a *multiscale decomposition* of an input image into layers at different scales of coarseness. In comparison with existing edge-preserving multiscale decompositions, our algorithm significantly increases the ability to distinguish high-contrast texture from a dense field of edges.

Our notion of detail inherently captures repetitive variation of intensity, which we term *oscillations*. Locally, the amplitudes of oscillations represent contrast while their spatial-frequencies represents fineness in scale. Our goal is to smooth fine-scale oscillations, or *detail*, regardless of their amplitudes (see Fig 6). We extract the locally finest-scale oscillations as detail using a single smoothing operation, and obtain a multiscale decomposition by progressive smoothing. During successive smoothing operations on the residual, we coarsen the scale at which extrema are detected.

Inspired by empirical mode decomposition and morphological image filters, we examine the *local extrema* of the input image to detect oscillations. Empirical decomposition does not preserve edges
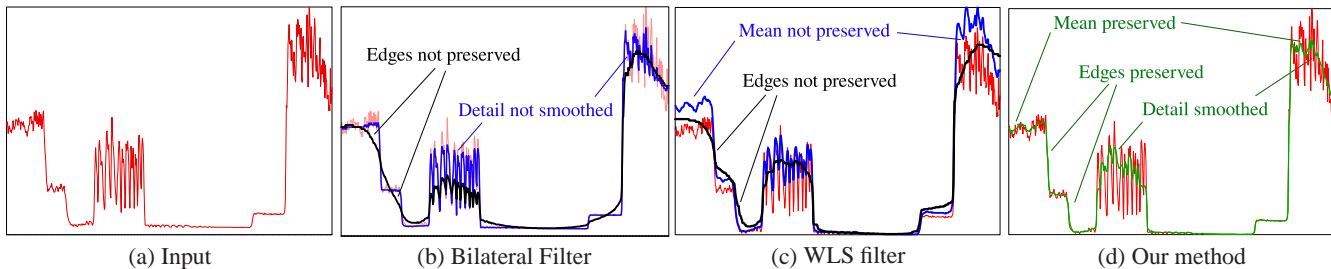
| (a) Input | (b) Bilateral Filter | (c) WLS filter | (d) Our method |

Figure 2: *Intensity plots along a scanline of an input image are shown with three filtered versions: (b) Bilateral filtering with a conservative (blue) and aggressive (black) range parameter values ; (c) Gradient-based edge preserving smoothing technique (WLS filtering [2008]) with larger (blue) and smaller (black) gradient preserving parameter values; (d) Our smoothing filter. While existing techniques (b) and (c) are effective in smoothing variation with small amplitude (blue), they necessarily blur edges (black) that have smaller magnitudes of gradients than the oscillations to be smoothed. Our smoothing algorithm smoothes large oscillations and strictly preserves edges (green), without the need for careful selection of parameter values.*

	Definition of Detail	Definition of Edge	Assumption
Bilateral Filter	Low contrast variation	Large intensity difference	Texture is low-contrast
WLS Filter	Low contrast variation	Large gradient	Magnitude of gradient is larger at edges
BEMD	Fine-scale spatial oscillations	No notion of edge	Means of oscillations are smooth
Our Algorithm	Fine-scale spatial oscillations	High variance in range values of neighboring local extrema	Detail is oscillations between local extrema

Figure 4: *Comparison of our approach with three existing techniques for image decomposition: Bilateral filtering [Fattal et al. 2007] , weighted least squares (WLS) filtering [Farbman et al. 2008] and bidimensional empirical mode decomposition (BEMD) [Huang 1998] .*

while morphological operations do not preserve shape [Serra and Vincent 1992]. We exploit information provided by local extrema about the oscillations in the image and preserve both— edges and shape. Our algorithm is based on two key observations: (1) Detail (even if high-contrast) is characterized by a large density of local extrema; (2) salient edges (even if low-contrast) are characterized by a large variation in their neighboring extremal values.

Using local extrema, rather than contrast, to characterize detail provides two important benefits. First, we make no *a priori* assumptions on the dynamic range of the input image or on the amplitude of the oscillations. Second, we obtain the local scale of oscillations independent of contrast. Progressive coarsening of the scale at which extrema are detected results in layers with oscillations at different scales. Also, by recursively removing detail, the degrees of coarseness in the multiscale decomposition are likely to capture the inherent superimposed scales of oscillation in the input image.

For simplicity, we describe our algorithm for an input grayscale image I. Similar to existing decomposition techniques, we perform the decomposition on the luminance channel for color images. We denote image-space coordinates (x, y) with boldface letters. Thus $I(\mathbf{p})$ is the intensity of the given grayscale image I at pixel $\mathbf{p}$.

2.1 Smoothing

We define *detail* as *oscillations between local minima and maxima* (see Fig 5). We extract detail by subtracting a smoothed image, that we call the *mean*, from the input. The smoothing algorithm uses the local extrema to detect oscillations at their finest scale, locally. By interpolating the minima and maxima independently, we construct two *extremal envelopes*, that sandwich the data, and propagate information about local oscillations to all pixels in the image. The average of the two interpolants, evaluated at each pixel, provides an estimate of the local mean about which the oscillations occur. To ensure that the mean respects edges in the input image, the interpolants need to be edge preserving in the traditional sense that they retain fidelity to the input at strong gradients.

Our smoothing algorithm consists of three steps: (1) Identification of local minima and local maxima of I; (2) Interpolation of the local minima and maxima to compute minimal and maximal extremal envelopes respectively; (3) computation of the smoothed mean M as the average of the extremal envelopes. Figure 5 illustrates the three steps of our smoothing algorithm by plotting 1D slices of the Barbara input image (red), its extrema, extremal envelopes (blue and magenta) and smoothed mean (black). The detail layer is obtained as $D = I - M$.

Extrema location: We use a simple test for locating image maxima. Pixel $\mathbf{p}$ is reported as a maxima (resp. minima) if *at most* $k - 1$ elements in in the $k \times k$ neighborhood around $\mathbf{p}$ are greater (resp. smaller) than the value at pixel $\mathbf{p}$. Oscillations whose maxima are detected by using a $k \times k$ kernel have wavelengths of at least $k/2$ pixels. Intuitively, using a large kernel overlooks the detection of fine oscillations. We start with $k = 3$ and increase the kernel size for multiscale smoothing, after extracting fine oscillations (see Sec. 2.2).

Extremal envelope construction: Given an image I and a set of pixels S (image local extrema), we compute an extremal envelope E using an interpolation technique that was proposed by Levin et al. [2004] for image colorization. In our context, we seek an interpolant E such that neighboring pixels $E(\mathbf{r})$ and $E(\mathbf{s})$ have similar values if $I(\mathbf{r})$ and $I(\mathbf{s})$ are similar. More formally, we minimize the functional

$$\sum_{\mathbf{r}} \left(E(\mathbf{r}) - \sum_{\mathbf{s} \in N(\mathbf{r})} w_{\mathbf{rs}} E(\mathbf{s}) \right)^2 \qquad (1)$$

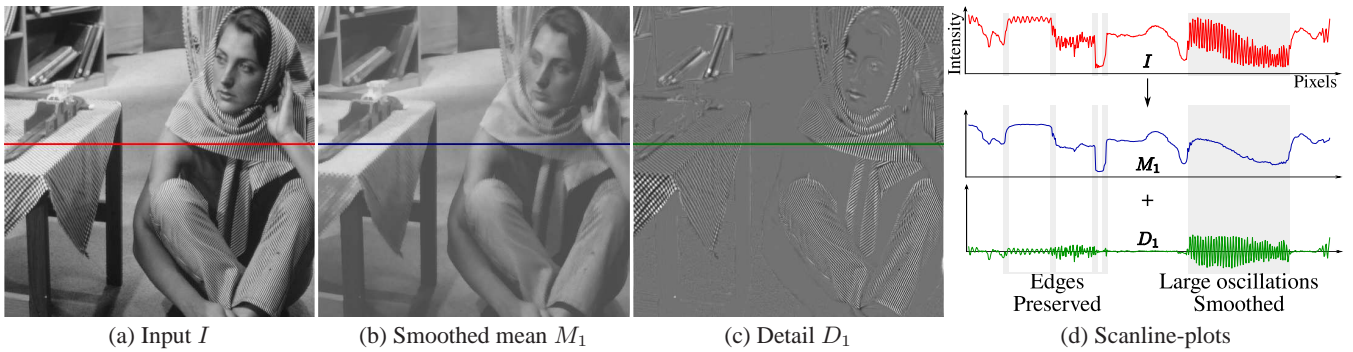| (a) Input I	(b) Smoothed mean M_1	(c) Detail D_1	(d) Scanline-plots

Figure 6: *Plots showing the input intensities (red) along a row and its separation into detail (green) and mean (blue) by our algorithm. Despite the large amplitude of some oscillations they are extracted as detail D_1, while single edges of lower amplitude are preserved in the smoothed mean M_1.*
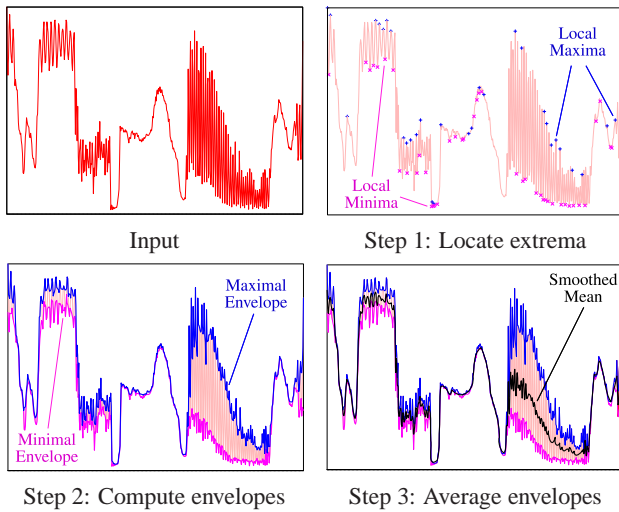

Figure 5: *The three steps of our smoothing algorithm illustrated with plots of intensity along the row shown in Figure 6.* Step 1: *We locate the local minima and maxima of the input (red). Note: The plot is along a row in the 2D input and extrema corresponding to some peaks seem to be missing since they lie on adjacent scanlines.* Step 2: *We compute the minimal (magenta) and maximal (blue) envelopes as edge-preserving interpolants through the minima and maxima respectively.* Step 3: *The smoothed mean (black) is computed as the average of the two envelopes.*

subject to the constraint

$$\forall\ \mathbf{p} \in S \quad E(\mathbf{p}) = I(\mathbf{p}).$$

$N(\mathbf{r})$ denotes the neighbors of $\mathbf{r}$, and weights

$$w_{\mathbf{rs}} \propto \exp\left(-\frac{(I(\mathbf{r}) - I(\mathbf{s}))^2}{2\sigma_r^2}\right) \quad (2)$$

are computed using the local variance σ_r^2 around $\mathbf{r}$. We adopt the approach of Levin et al. [2004] and minimize the quadratic functional using their weighted least squares formulation, which reduces to solving a sparse linear system with $N(\mathbf{r})$ defined as a 3×3 local neighborhood.

Smoothed mean: Performing the envelope construction independently on the minima and maxima of the image yields the minimal and maximal envelopes respectively. The smoothed mean image is computed as the average of these two envelopes (see Fig 5).

2.2 Multiscale decomposition

A single smoothing operation of I yields a detail image, D_1, that contains the finest-scale local oscillations and a mean, M_1, that represents a coarser trend. We obtain a multiscale decomposition of the input image by recursively extracting a number of detail layers from the mean. After n recursive smoothing operations, we obtain detail images $D_1, D_2, ..., D_n$ at increasing scales of coarseness and a residual mean image:

$$I(\mathbf{p}) = \sum_{i=0}^{n} D_i(\mathbf{p}) + M_n(\mathbf{p}). \quad (3)$$

Choosing $k = 3$ as the size of the extrema-location kernel (see Sec. 2.1) for the first smoothing step of I results in a detail D_1 that captures oscillations of frequency up to $3/2\,\mathrm{pixel}^{-1}$. By increasing k, we effectively capture coarser oscillations while recursively smoothing M_1. Progressively increasing k through each recursive smoothing causes the different detail layers to contain increasingly coarse oscillations. In our experiments we found that the algorithm was not sensitive to the factor by which k was increased. For all the results in the paper we increased k by a constant value of eight, between iterations. Figure 7(d) visualizes the extrema of I, M_1 and M_2. For compact visualization, the three sets of extrema are shown in different vertical regions of the image.

2.3 Discussion

Effects of noise: For noisy input images, our algorithm effectively separates the noise if the scale of the noise does not match the scale of features in the input image. We repeated an experiment performed by Farbman et al. [2008], on a greyscale image with several step-edges of varying magnitude that was polluted with noise at two scales. Our decomposition algorithm effectively recovers the noise at different scales (see Fig. 8).

Edge preservation: Current edge-preserving image decompositions use local contrast to define edges. On the other hand, we define edges as regions where the *variation in the values of the neighboring extrema* is large. Our smoothing filter preserves edges because the extremal envelopes implicitly maintain fidelity to the data at pixels where the variation in the range values of the nearby extrema is large. Regions with large-amplitude, oscillations are smoothed effectively since the local extrema have similar values.
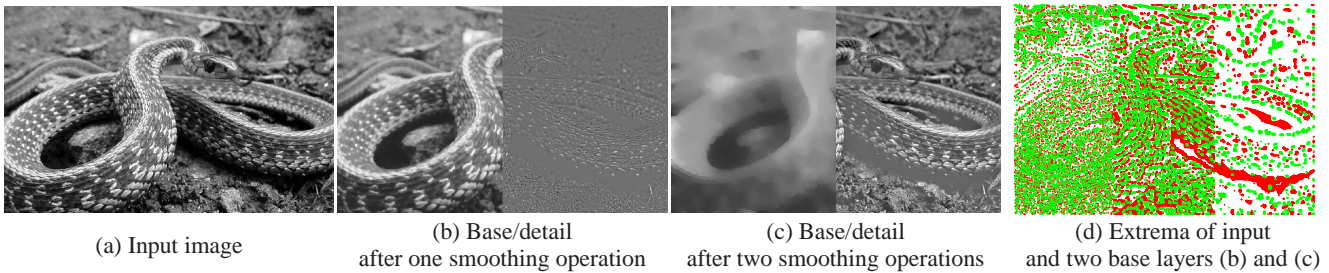
4

| (a) Input image | (b) Base/detail after one smoothing operation | (c) Base/detail after two smoothing operations | (d) Extrema of input and two base layers (b) and (c) |

Figure 7: *Our multiscale decomposition extracts features based on their spatial scale. An input image is shown along with its three-layer decomposition. The local extrema of the input image, the base layer in (b) and the base layer in (c) are shown as three abutting vertical regions in (d).*
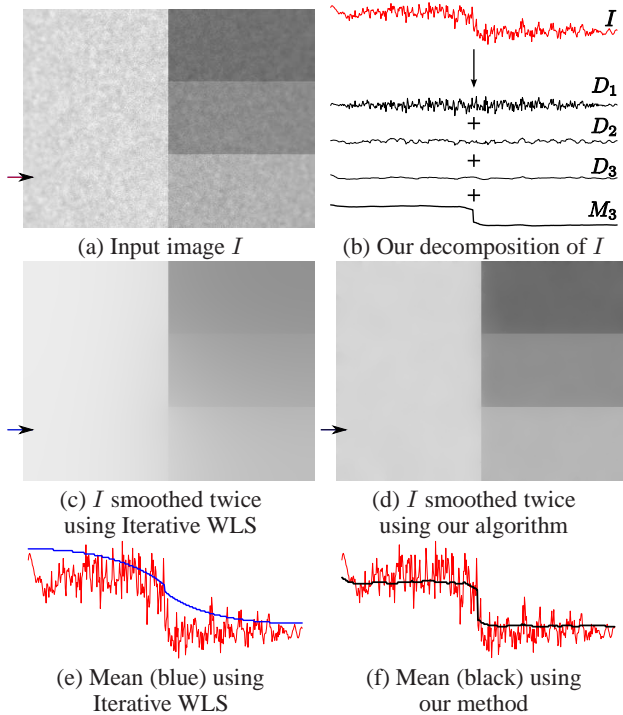

(a) Input image I	(b) Our decomposition of I
(c) I smoothed twice using Iterative WLS	(d) I smoothed twice using our algorithm
(e) Mean (blue) using Iterative WLS	(f) Mean (black) using our method

Figure 8: *Results of applying our algorithm on a noisy image (courtesy of Farbman et al. [2008]). (a) The input image I is a piecewise constant image containing several step-edges of different magnitudes, to which noise was added at different scales. Our smoothing algorithm produces a better estimate of the mean while effectively extracting detail at multiple scales. (b) The result of our decomposition on a single row. (c) The result of smoothing I using iterative WLS [Farbman et al. 2008]. (d) The result of smoothing I using our algorithm. (e) A plot of the smoothed result (blue) using WLS filtering, along with the input (red). (f) A plot of the our smoothed result (black) with the input (red).*

Robustness to image scaling: Performing the decomposition of a scaled version of an image provides consistent results if the window used for extrema detection is scaled accordingly. The size of the kernel used in our extrema detection determines the largest frequency of oscillations that can be extracted as detail. To maintain consistency between decompositions of scaled versions of the input image it suffices to simply scale the the kernels by the same factor.

Sparse extrema: When the density of local extrema is very low, the

interpolation [Levin et al. 2004] can become unstable. However, a low extremal density indicates that the underlying function is very smooth. Introducing artificial interpolation constraints (extrema) in smooth regions makes the interpolation stable. In practice, we insert artificial extrema in regions of the image that contain no extrema and are larger than a given threshold size (50×50 pixels).

Smoothing by contrast reduction: In traditional empirical mode decomposition [Huang 1998] of smooth 1D data, smooth interpolation schemes are used to construct the extremal envelopes. We use an edge-preserving interpolation scheme so that the smoothed mean preserves isolated discontinuities. The tendency of the interpolant to preserve large gradients may result in incomplete smoothing of oscillations in a single iteration. However, a combination of increasing the window size for extrema-location and performing the decomposition in the log-domain make this effect almost imperceivable. Another solution is to repeat each smoothing step (keeping k fixed) until the detail is completely extracted.

Features at boundaries of textured regions: large-amplitude oscillations that occur at the boundaries of textured regions are indistinguishable from edges. Figure 9 illustrates an example where, despite the high contrast, the spotted pattern on the hat is smoothed effectively while subtler shading is preserved on the coarse scale. However, the bright spots at the boundary with the ribbon are mistaken to be part of the ribbon. Handling such cases would require semantic information such as from an explicit pattern matching algorithm.

3 Results

We tested our smoothing and decomposition algorithms on a variety of images. On average, a four-layer decomposition of 1024×768 images took about 30 seconds using a naïve solver for computing the extremal envelopes. Using a simple multigrid solver, we were able to achieve a speedup of about 1.5. To locate extrema, we use a 3×3 kernel for the finest detail and progressively enlarge the kernel by a constant value (8) through the recursion for coarser layers.

3.1 Comparison

We wish to stress the difference in philosophies between current algorithms and our approach. Our novel definition of detail, as repetitive oscillatory features between local extrema, produces fundamentally different decompositions from existing solutions that interpret large gradients as edges to be preserved. The differences are primarily with coarse-scale features that have low contrast and fine-scale features that are highly contrasted.

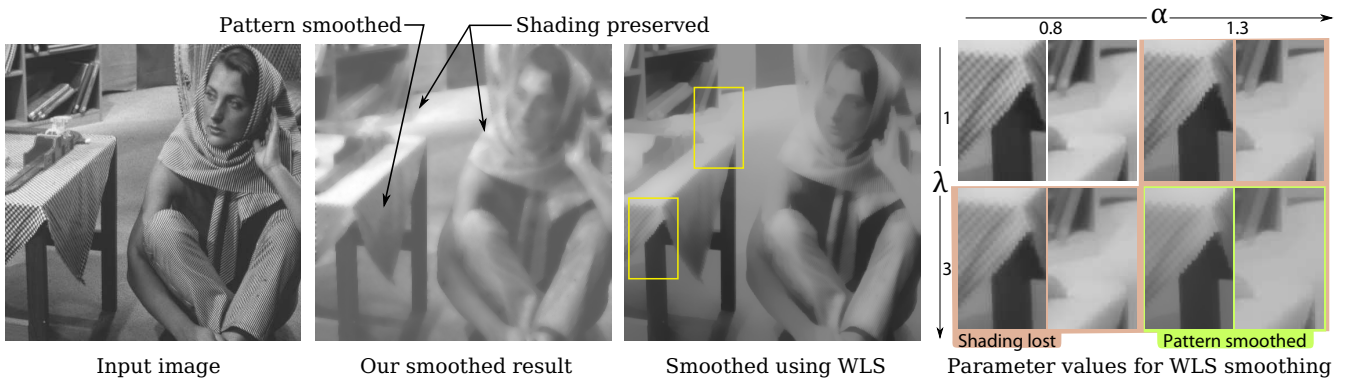Techniques, that extract low contrast features as detail, typically

Figure 10: *The Barbara input image along with results of smoothing with the WLS filter [Farbman et al. 2008] using various combinations of the input parameters. Zooming into insets with contrasted texture and subtle shading, we see that gradient-based techniques are unable to preserve subtle, coarse features while smoothing fine, well-contrasted texture. Our method preserves subtle shading and effectively smooths the texture.*

(a) Input image   (b) Smoothed image

Figure 9: *Failure case: Although the high-contrast, spotted pattern on the hat is smoothed effectively while retaining subtler shading information, parts of this pattern on the boundary with the ribbon are indistinguishable from the ribbon. Although our definition of detail does not inherently disambiguate edges from partial oscillations of similar amplitude at boundaries, this is an extreme example. Handling such cases would require semantic information such as from an explicit pattern matching algorithm.*

demonstrate their utility using images where the low contrast detail also tends to be fine-scale. On such images, despite the difference in philosophies, our results are quite similar since fine-scale features extracted by our technique as detail also happen to be of low contrast. For example, using the flower example of Farbman et al. [Farbman et al. 2008] we achieve similar results (Fig. 11) since the details on the flower petals are fine-scale and of lower contrast than at the boundaries. In this paper, we focus on cases that produce different decompositions from gradient-based approaches.

Figure 15 compares the results of our technique with existing decomposition schemes. One key difference is that our decomposition extracts, earlier, fine-scale features (such as the pebbles towards the bottom of the image) as detail, while existing schemes extract low-contrast features (such as the large clouds) earlier as detail.

Figure 10 shows an example where the input contains texture that is more contrasted than some edges. Using a purely gradient de-

pendent approach, smoothing the oscillation necessarily smoothes low-contrast edges (see also Fig. 3). Also, current decompositions can involve non-intuitive manipulation of input parameters across different images. In comparison, our technique is simple, smoothes texture, respects soft, single edges, preserves subtle shading and consistently smoothes a variety of images with widely different contrasts.

3.2 Applications

Multiscale decompositions of images, into layers of varying contrast, have been used in several applications including equalization and image abstraction [Farbman et al. 2008; Lischinski et al. 2006; Fattal et al. 2007]. In addition to these, we present applications that exploit a key property of our decomposition— the extracted layers correspond to superposed oscillations of increasing coarseness. We apply our decomposition to enhance detail (image equalization) and to remove detail (estimating tone from cross-hatched images, separating texture from illumination, illumination transfer).

Hatch to tone: Few techniques are able to recover tone from images with hatching or stippling, while preserving edges. The difficulty lies in retaining edges depicted by these techniques while smoothing high-contrast variation. Smoothing filters like the bilateral filter or weighted least squares filter are not very useful in this context. Figure 12 shows the residual from running three iterations of our smoothing algorithm on a cross-hatched input image. We smooth fine-scale oscillations, ideally to their flat means, earlier in the process. However, in the case of non-homogeneous, high-contrast oscillations, the edge preserving nature of the non-linear extremal interpolants causes the contrast of the oscillations to be reduced considerably but not completely. Consequently the computed mean tends to contain residual oscillations that are grayscale. The amplitude of these residual oscillations depends on its original wavelength; fine oscillations leave weaker residuals than coarse ones. Over multiple iterations of such smoothing applied on binary (or highly contrasted) hatched images, the complex interplay between homogeneity of oscillations in 2D and grayscale residuals from previous iterations tends to result in a smoothed image where the tone at each pixel is directly related to the frequency of local oscillations. While we smooth variation, the edges of variations are well preserved. We compare our solution with a median filter. The problem with the latter is that, using a small kernel size, tone is not recovered at a coarse scale and Increasing the kernel size wipes out thin features like outlines. Another drawback of the median filter is

6

Figure 11: *Fine-scale enhancement of the input image (left) using WLS [Farbman et al. 2008] (middle) and our technique (right) provides similar results with subtle differences since the detail on the petals are of, both, low contrast and fine scale. The WLS method fails to enhance fine-scale detail that are high-contrast such as the serrations on the leaves in the background and specularities on the small leaves on the right. In addition, coarser features such as the subtle discoloration on the defocused top-left portion of the image are enhanced as detail by the WLS method simply because they are low contrast.*

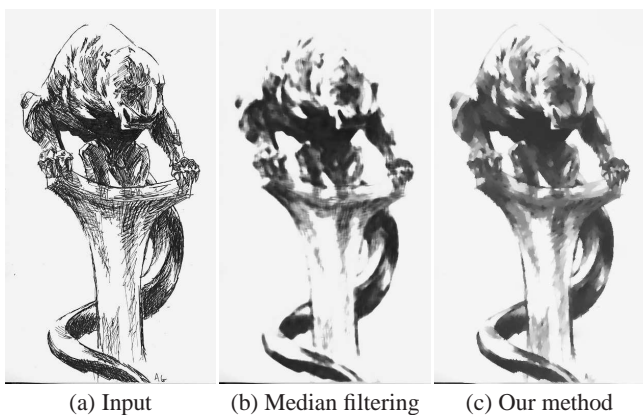that the filter only selects pixel levels that are present in the input image.

(a) Input   (b) Median filtering   (c) Our method

Figure 12: *(b) Applying a median filter has two disadvantages: Choosing a large kernel size washes out thin edges while choosing a small kernel size does not smooth the hatched pattern; also the median filter simply selects one of the existing grey levels and cannot produce intermediate shades of grey. (c) The residual after three iterations of smoothing using our algorithm yields a good estimate of the tone while preserving the edges of hatched regions.*

Separating fine texture and coarse shading: We are able to separate fine texture from shading, provided the oscillations of the texture and shading are of different scales. Although we make the same assumption as Oh et al. [2001] that illumination information is "lower frequency" than texture, we do not make any assumptions on the contrast of the texture. Since Oh et al. use the bilateral filter, they are prone to the additional assumption that the contrast of the texture and shading are vastly different. We demonstrate the effectiveness of our algorithm by retexturing an image containing high-contrast texture, while retaining shading on the newly painted texture (see Fig. 13). We achieve this by transferring the coarsened luminance of the input image onto its edited version.

Image equalization: The layers from our decomposition can be seen as an adaptive basis that sum to the input. By considering different linear combination of these layers, we show that detail at different scales can be exaggerated. In practice, since we manipulate the log-luminance channel, we perform the linear combinations in log space. Current equalization techniques define detail as low contrast. Instead, we are able to control relative contrasts of fea-
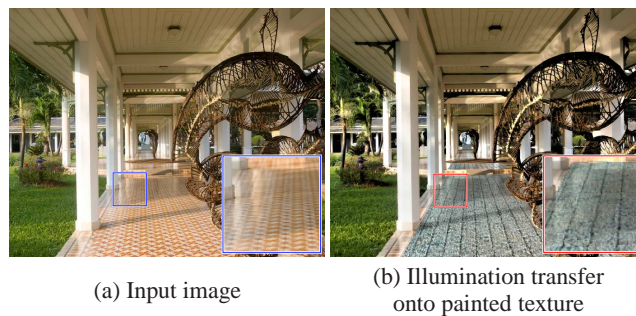

(a) Input image   (b) Illumination transfer onto painted texture

Figure 13: *Our edge preserving decomposition separates an input image into layers containing detail at different scales. (a) The tiled texture on the floor is finer than illumination effects such as glossy reflections and shadows. (b) The coarse illumination information is extracted from (a) and combined with the fine texture information extracted from (b) to preserve shadows and subtle effects such as glossy reflections of pillars on the newly painted texture (inset).*

tures based on their scales (see Fig.1). More examples of image equalization are presented in the video.

High dynamic range (HDR) images: Although filters that extract detail based on contrast (WLS and bilateral filters) are more appropriate tools for tone-mapping, in practice, we find that our equalizations produce reasonable results (see Fig. 14). An advantage of our method is intuitive and consistent parameter values across different images. However, since we filter based on scale and not contrast, specialized techniques may be preferable for input where the HDR content is spread across significantly different spatial scales.

4 Conclusion

We have presented a novel definition for image detail as oscillations between local minima and maxima. While existing decomposition algorithms extract detail based on a notion of contrast, our definition of detail captures the scale of spatial oscillations, locally.

Building on our definition of detail, we proposed a simple algorithm to smooth an input image. By recursively performing the smoothing with extrema detection at multiple scales, we performed a decomposition of the input image into multiple-scale layers of detail and a coarse residual. Our algorithm smoothes high-contrast texture while preserving salient edges. Finally, we exploited this ability by applying our decomposition in a variety of applications.
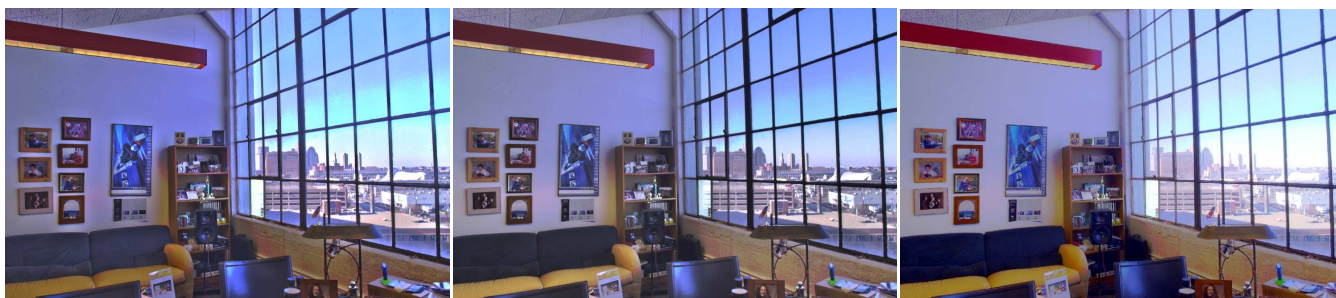
(a) Tone-mapped using the bilateral filter   (b) Tone-mapped using the WLS filter   (c) Our equalized result

Figure 14: *Comparison of our equalization against tone-mapping operators on an example high dynamic range (HDR) input image. (a) and (b) have been directly taken from [Durand and Dorsey 2002] and [Farbman et al. 2008] respectively. (c) is obtained using our 2-layer equalization where the base layer is scaled to half and recombined with the detail. Although our notion of detail is based on spatial scale and not contrast, our equalization can be used to achieve basic tone-mapping by scaling down the layer(s) with HDR content.*

Acknowledgements

References

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

BURT, P. J., AND ADELSON, E. H. 1983. The laplacian pyramid as a compact image code. *IEEE Trans. on Communications COM-31,4*, 532–540.

CHEN, J., PARIS, S., AND DURAND, F. 2007. Real-time edge-aware image processing with the bilateral grid. *ACM Transactions on Graphics*, 103.

CHOUDHURY, P., AND TUMBLIN, J. 2005. The trilateral filter for high contrast images and meshes. In *SIGGRAPH '05: ACM SIGGRAPH 2005 Courses*, ACM, New York, NY, USA, 5.

DAMERVAL, C., MEIGNEN, S., AND PERRIER, V. 2005. A fast algorithm for bidimensional emd. *Signal Processing Letters, IEEE 12*, 10 (Oct.), 701–704.

DURAND, F., AND DORSEY, J. 2002. Fast bilateral filtering for the display of high-dynamic-range images. In *ACM Transactions on Graphics: SIGGRAPH '02*, ACM Press, New York, NY, USA, 257–266.

FARBMAN, Z., FATTAL, R., LISCHINSKI, D., AND SZELISKI, R. 2008. Edge-preserving decompositions for multi-scale tone and detail manipulation. *ACM Transactions on Graphics*, 67.

FATTAL, R., AGRAWALA, M., AND RUSINKIEWICZ, S. 2007. Multiscale shape and detail enhancement from multi-light image collections. *ACM Transactions on Graphics*, 51.

HUANG. 1998. The empirical mode decomposition and the hilbert spectrum for nonlinear and non-stationary time series analysis. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences 454*, 1971 (March), 903–995.

LAGENDIJK, R. L., BIEMOND, J., AND BOEKEE, D. E. 1988. Regularized iterative image restoration with ringing reduction. *IEEE Trans. on Signal Processing (Acoustics, Speech, and Signal Processing) 36*, 12, 1874–1888.

LEVIN, A., LISCHINSKI, D., AND WEISS, Y. 2004. Colorization using optimization. *ACM Transactions on Graphics 23*, 689–694.

LI, H., YANG, L., AND HUANG, D. 2005. The study of the intermittency test filtering character of hilbert-huang transform. *Mathematics and Computers in Simulation 70*, 1, 22–32.

LISCHINSKI, D., FARBMAN, Z., UYTTENDAELE, M., AND SZELISKI, R. 2006. Interactive local adjustment of tonal values. *ACM Transactions on Graphics 25*, 3, 646–653.

LIU, Z., AND PENG, S. 2005. Boundary processing of bidimensional emd using texture synthesis. *Signal Processing Letters, IEEE 12*, 1 (Jan.), 33–36.

NUNES, J., NIANG, O., BOUAOUNE, Y., DELECHELLE, E., AND BUNEL, P. 2003. Texture analysis based on the bidimensional empirical mode decomposition with gray-level co-occurrence models. *Signal Processing and Its Applications, 2003. Proceedings. 2* (July), 633–635 vol.2.

OH, B. M., CHEN, M., DORSEY, J., AND DURAND, F. 2001. Image-based modeling and photo editing. In *Proceedings of SIGGRAPH 2001*, ACM, NY, USA, 433–442.

PATTANAIK, S. N., FAIRCHILD, M., FERWERDA, J., AND GREENBERG, D. P., 1998. Multiscale model of adaptation, spatial vision and color appearance.

RAHMAN, Z. U., AND WOODELL, G. A. 1997. A multi-scale retinex for bridging the gap between color images and the human observation of scenes. In *IEEE Trans. on Image Processing: Special Issue on Color Processing 6(7*, 965–976.

SERRA, J., AND VINCENT, L. 1992. An overview of morphological filtering. In *Circuits, Systems and Signal Processing*, 47–108.

TOMASI, C., AND MANDUCHI, R. 1998. Bilateral filtering for gray and color images. In *In Proc. of the Sixth International Conference on Computer Vision, Bombay, India, January 1998*.

TUMBLIN, J., AND TURK, G. 1999. Lcis: a boundary hierarchy for detail-preserving contrast reduction. In *Proceedings of SIGGRAPH '99*, ACM Press/Addison-Wesley Publishing Co., NY, USA, 83–90.
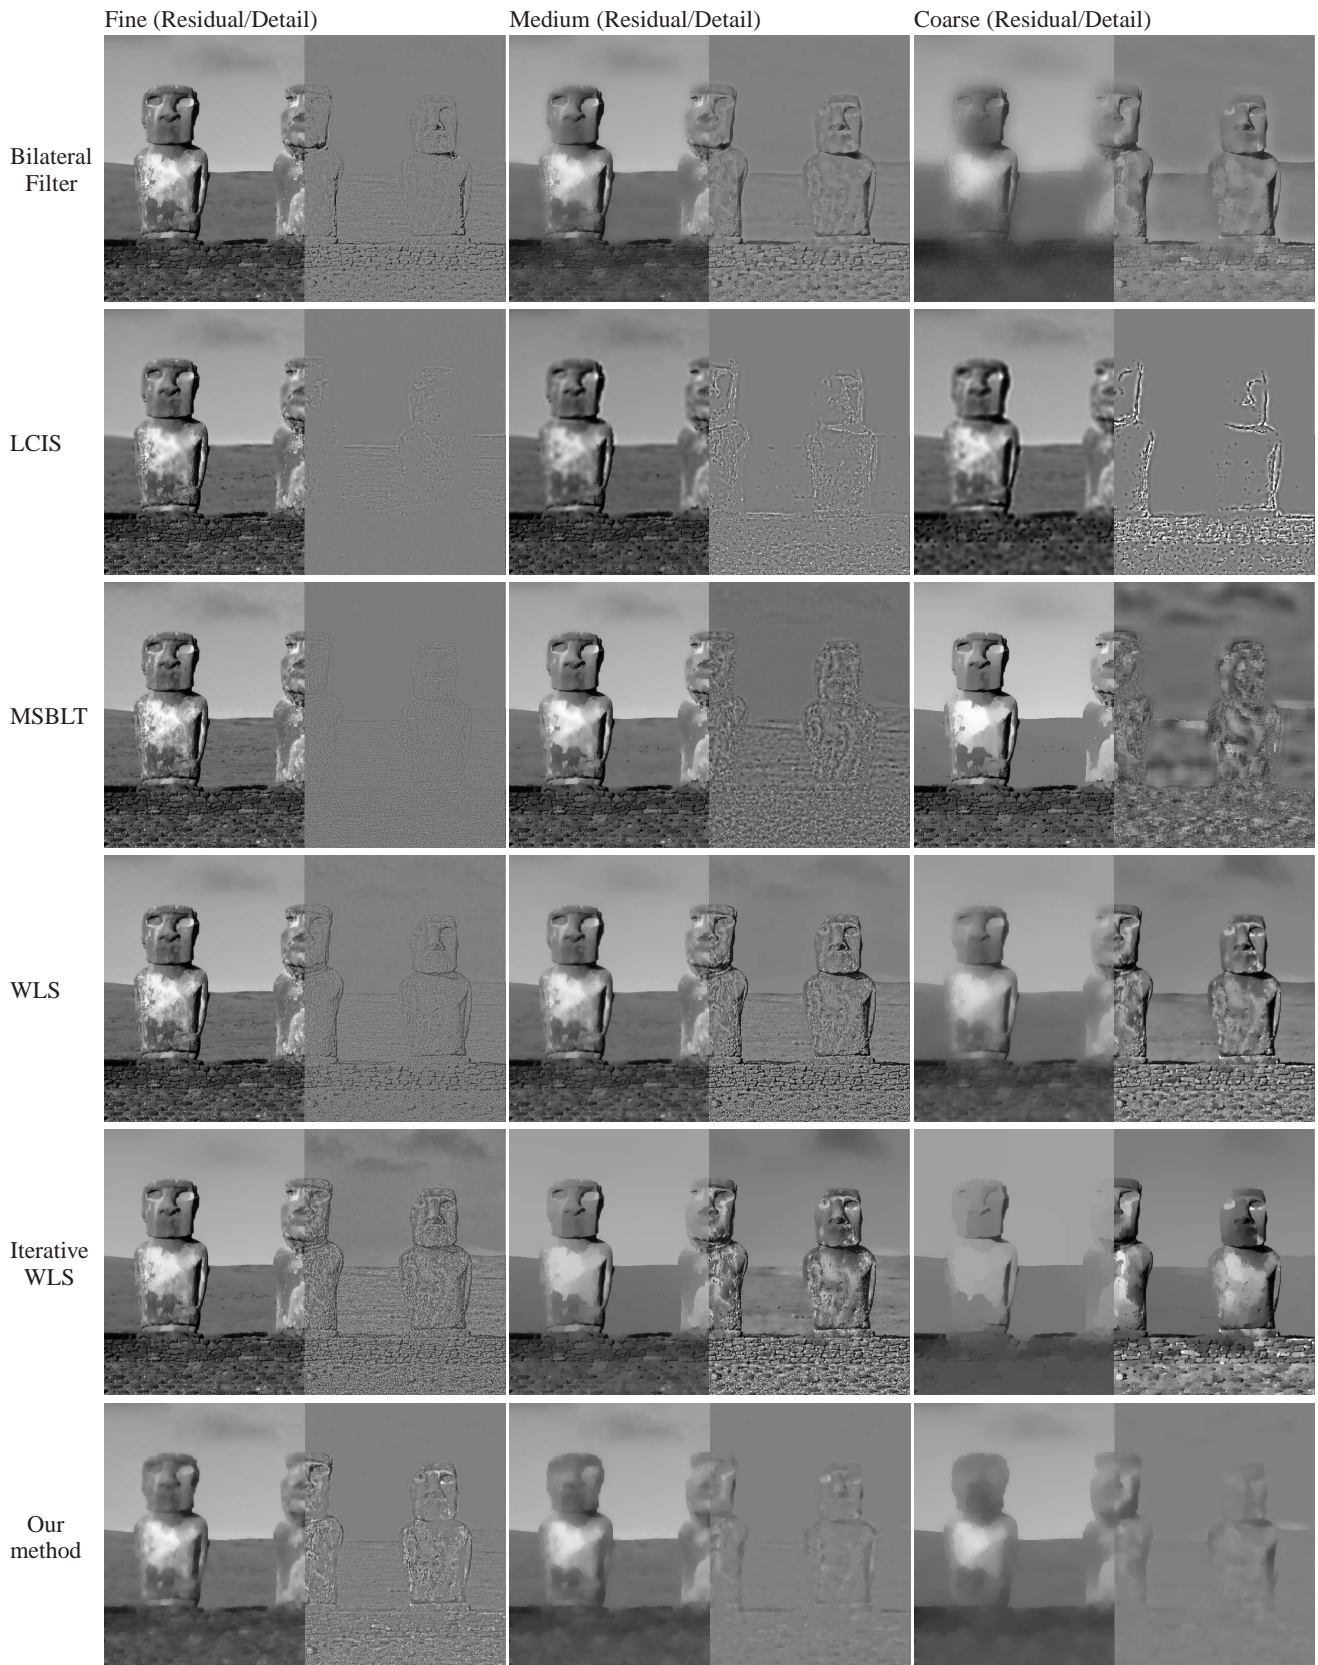
Fine (Residual/Detail)   Medium (Residual/Detail)   Coarse (Residual/Detail)

Bilateral Filter

LCIS

MSBLT

WLS

Iterative WLS

Our method

Figure 15: *Comparison of our results with existing approaches: Bilateral filtering [Chen et al. 2007], MSBLT [Fattal et al. 2007], LCIS [Tumblin and Turk 1999], WLS, iterative WLS [Farbman et al. 2008]. Our smoothing extracts features based on spatial scale while other methods smooth low-contrast features first. Using our decomposition, the pebbles and stones towards the bottom of the image are extracted as fine- and medium-scale detail respectively, even though they are well contrasted. On the other hand, despite their low contrast, the clouds are not extracted as detail due to their coarse scale. The comparison images have been directly taken from [Farbman et al. 2008].*