An Extended Exposure Fusion and its Application to Single Image Contrast Enhancement

Charles Hessel Jean-Michel Morel
Université Paris-Saclay, ENS Paris-Saclay, CNRS, Centre de mathématiques et de leurs applications, 94235 Cachan, France
{charles.hessel, morel}@ens-paris-saclay.fr

Abstract

Exposure Fusion is a high dynamic range imaging technique fusing a bracketed exposure sequence into a high quality image. In this paper, we provide a refined version resolving its out-of-range artifact and its low-frequency halo. It improves on the original Exposure Fusion by augmenting contrast in all image parts. Furthermore, we extend this algorithm to single exposure images, thereby turning it into a competitive contrast enhancement operator. To do so, bracketed images are first simulated from a single input image and then fused by the new version of Exposure Fusion. The resulting algorithm competes with state of the art image enhancement methods.

1. Introduction

Introduced in 2007 by Mertens *et al*. [28, 29] Exposure Fusion (EF) is a high dynamic range imaging technique fusing a bracketed exposure sequence into a high quality image. Contrarily to most HDR imaging methods, Exposure Fusion does not create an intermediate HDR image but directly constructs the final LDR image by seamlessly fusing the best regions of the input sequence, using the Laplacian pyramid. Since its publication, this method has received considerable attention, being both effective and efficient.

In this paper, we propose a simple solution to two notorious artifacts of exposure fusion, namely the out-of-range artifact and the low-frequency halo. We call the improved algorithm extended exposure fusion (EEF), and show that not only its fusion results are artefact free, but also that it can be used to produce directly images with increased local contrast. Visual comparisons demonstrate that the proposed method outperforms state-of-the-art fusion techniques. Further, we derive a single image contrast enhancement technique from this corrected exposure fusion. To this aim, a bracketed exposure sequence is simulated from the single input and then fused. The enhancement method inherits all the qualities of extended exposure fusion.

Our plans is as follows. After a brief description of EF in Section 3, we explain and solve in Section 4 the out-of-range artifact and the low-frequency halo. This section defines an *extended exposure fusion* (EEF), which fuses real bracketed exposure sequence without artifacts. Results are compared against state-of-the-art methods. We also present an application to the fusion of two large-exposure-ratio images. The EEF method is then adapted to the enhancement of a single exposure in Section 5. We thus obtain a single image contrast enhancement method that we call *simulated exposure fusion* (SEF). A comparison of SEF with state-of-the-art methods is carried out in Section 6.

2. Related Work

The classic HDR imaging pipeline consists in the construction an HDR image by fusion of a bracketed exposure sequence followed by the application of a tone-mapping operator [37]. Since the publication of the Mertens *et al*. paper [28, 29] however, exposure fusion methods have emerged as an attractive alternative to this pipeline, as they avoid a camera response curve calibration and artifact-prone tone-mapping steps. EF is not the first method of the kind though [4, 8]. In the last decade, numerous exposure fusion methods have been proposed that changed either the quality metrics or the fusion method, sometimes both. Yet they all keep the same construction, namely the creation of weight maps using quality metrics, and a weighted blending.

Regarding the quality measures, Mertens *et al*. [28] used contrast, namely the amplitude of the Laplacian, color saturation and well-exposedness. The magnitude of the Laplacian had been used before [2], the color saturation was reused in [39, 41] and some other metrics were proposed, *e.g*. the entropy [8, 12], saliency [21], or the amount of detail after a base+detail decomposition using an edge-preserving smoothing filter [36, 39]. SSIM-related [45] image quality metrics have also been used in [25, 27].

As for the fusion strategy, numerous variations have been explored. Fusion can be carried out through global optimization [10, 16, 35, 39], in which case the quality is directly measured in the final image. Other fusion techniques are single-scale [12, 20, 36, 39–41], or pyramid-based [2, 4, 19, 29]. The former generally require to refine the weighting maps, *e.g*. using the (weighted) guided filter [21–23]. Kou *et al*. [17] proposed a edge-aware smoothing pyramid. The gradient domain was used in [9,14,43,46]. Base+detail decomposition was also used, either to compute the blending weights [36], or to enhance the local contrast: Wang *et al*. [44] enhance the local contrast of the input sequence using the Local Laplacian Filter [31, 32] before fusing the images, whereas Li *et al*. [22] add a detail-enhancement step after fusion. Mixing these ingredients, Li *et al*. [49] proposed an optimization-based method in the gradient domain, where a detail layer is computed and added to the fused image to enhance its local contrast.

None of these works solved the initial issues of the Mertens *et al*. method, though described by the authors themselves [29]. Kou *et al*. [17] succeeded, however, to circumvent the out-of-range artifact by reducing the depth of the pyramid. But doing this increases the low-frequency halo. To avoid it, the authors introduced an edge-aware smoothing pyramid. Our approach has two major differences: first, it is simpler and does not rely on any additional smoothing step, but on mere remapping functions. Our method tackles the artifact at its root. Moreover, we will see in Section 4 that [17] eventually creates some halos. Second, rather than reducing the depth of the pyramid, we increase it. This is better, since a deeper pyramid gives more natural results, with preserved relative brightness.

Image contrast enhancement techniques belong to the wider class of tone-mapping operators. While this term generally refers to methods that map values of a high dynamic range (HDR) image into the smaller range of a low dynamic range (LDR) image, the very same methods can generally be adapted to LDR images and improve their brightness and contrast. *Global* operators like histogram equalization and gamma-correction apply the same correction to pixels with the same color, whereas *local* operators map the values depending on the local content. Among the many local methods introduced for LDR images we can mention ACE [5,6], LCC [7, 24, 38], LLF [1, 31]. Contrast enhancement techniques take inspiration in Land's retinex method [18] which modeled the human color perception system. Its most popular implementation is arguably MSR [15, 33].

Given the success of the exposure fusion methods, the number of papers that suggested to simulate a sequence of images from a single one, and then fuse it, is surprisingly small. Lee *et al*. [19] proposed to generate virtual exposure images using a function imitating the F-stop concept. They classify the input image in three classes: dim, well-exposed and bright, and adapt the well-exposedness measure. A multiscale fusion is then carried out in the discrete wavelet transform domain. This method is heavier than ours: while we generally uses only five or less images thanks to our histogram-based generation process, they usually need twice more. Besides, their method can reduce local contrast. Our method secures local contrast preservation or enhancement everywhere in the image. In their paper on burst photography for mobile cameras, Hasinoff *et al*. [11] also used EF with simulated images, as a way to perform tone-mapping at a very low cost. They simulate two images, one with short and one with long exposure, from an HDR image obtained from the fusion of a burst of images with constant exposure time. Lastly, in their 2017 paper on image fusion, Li *et al*. [22] proposed to simulate images and fused with their modified version of the Mertens *et al*. method [28]. They aim at enhancing backlit images by fusing three images with increased contrast in the dark parts, leaving bright regions untouched. Their results suffer from strong out-of-range artifacts, causing loss of contrast in the bright regions (see Section 6).

Our improvement of the Mertens *et al*. Exposure Fusion cancels out artifacts and increases local contrast. This is relevant, considering that EF, despite its numerous derivations, was still judged in 2015 the best available [26].

3. Exposure Fusion

Exposure fusion [28, 29] first measures the perceptual quality of the input sequence's images using three pixel-wise metrics: contrast C, saturation S and well-exposedness E. We will denote in the following by u the image, by $\mathbf{x}=(x_1,x_2)$ the position of the pixel, by c the color channel, and by k the position of the image in the sequence. The *contrast metric* uses the absolute value of a discrete Laplacian filter applied to the grayscale version of the image:

$$C_k(\mathbf{x}) = \left|\left(\tfrac{1}{3}\sum_{c=1}^{3} u_{c,k}\right) * K_{\text{Laplacian}}\right|(\mathbf{x}). \tag{1}$$

The authors use for the Laplacian kernel $K_{\text{Laplacian}}$ the sum of differences over the four nearest neighbors. The *saturation metric* is the standard-deviation of the pixel's color,

$$S_k(\mathbf{x}) = \sqrt{\tfrac{1}{3}\sum_{c'=1}^{3}\left(u_{c',k}(\mathbf{x}) - \tfrac{1}{3}\sum_{c=1}^{3} u_{c,k}(\mathbf{x})\right)^2}. \tag{2}$$

Finally, the *well-exposedness* measures how close each pixel's color is to the median value 0.5 using a Gauss curve:

$$E_k(\mathbf{x}) = \prod_{c=1}^{3} \exp\frac{-(u_{c,k}(\mathbf{x})-0.5)^2}{2\sigma^2}, \tag{3}$$

with $\sigma=0.2$. We follow the convention that the images' displayable range is $[0,1]$. The quality measure of each pixel is finally obtained as a product of these three metrics:

$$w_k(\mathbf{x}) = C_k(\mathbf{x})^{\omega_c} \times S_k(\mathbf{x})^{\omega_s} \times E_k(\mathbf{x})^{\omega_e}. \tag{4}$$

Power functions with parameters ω_c, ω_s and ω_e are added to allow the user to balance the importance of the different metrics. In this paper, unless notified otherwise, we keep

Figure 1. Out-of-range artifact. The bracketed sequence (left) is fused into a single output with EF (center, right). The result has out-of-range values, *e.g.* in the windows' bright parts: its dynamic range is $1.75\times$ larger than the input's one. EF clips the fused image (center), thus loses information. The normalized result (right) has a drastically reduced contrast compared to the input images.

these parameters equal to one. For the blending process, the resulting weights need to be normalized as

$$\widehat{w}_k(\mathbf{x}) = w_k(\mathbf{x}) \left(\sum_{k'=1}^{N} w_k(\mathbf{x}) \right)^{-1}. \tag{5}$$

The input images are then fused according to their normalized weight maps using the Ogden *et al.* multiscale fusion [30]. This technique builds the Laplacian pyramid [3] of the output image by blending the Laplacian pyramids of the input images according to the Gaussian pyramid of the weight maps. The fused image is obtained by collapsing the constructed pyramid. We will denote $L\{u\}$ the Laplacian pyramid of u, $G\{w\}_k$ the Gaussian pyramid of the weights w_k, and l the scale. The blending operation is then:

$$L\{v\}^l(\mathbf{x}) = \sum_{k=1}^{N} G\{\widehat{w}\}_k^l(\mathbf{x}) L\{u\}_k^l(\mathbf{x}). \tag{6}$$

Algorithm 1 describes the whole process, from the quality measurements to the multiscale fusion.

Algorithm 1: Exposure Fusion (EF)

input : sequence of images u, parameters $\omega_s, \omega_c, \omega_e$
output : fused image v
foreach *image at index* $k \in \{1, 2, ..., N\}$ **do**
  Compute metrics C, S, E // Eq. (1), (2), (3)
  Compute weight map w_k // Equation (4)
Compute the normalized weights $\widehat{w}_k$ // Equation (5)
Initialize output pyramid $L\{v\}$ with zeros
foreach *image at index* $k \in \{1, 2, ..., N\}$ **do**
  Compute $G\{\widehat{w}\}_k$ // Gaussian pyramid of $\widehat{w}_k$
  Compute $L\{u\}_k$ // Laplacian pyramid of u_k
  foreach *coefficient of the pyramid* **do** // Eq. (6)
    $L\{v\}^l(\mathbf{x}) \leftarrow L\{v\}^l(\mathbf{x}) + G\{\widehat{w}\}_k^l(\mathbf{x}) L\{u\}_k^l(\mathbf{x})$
$v \leftarrow$ collapse Laplacian pyramid $L\{v\}$

4. Correcting exposure fusion

Fusion methods should preserve the relevant information from all input images. In Exposure Fusion, the desirable image contains the well-exposed areas from the input bracketed sequence (and without distortions!). However EF sometimes fails at constructing such an image. The authors themselves described two artifacts affecting their results [29]: first, an expansion of the fused image's dynamic range with respect to the inputs (out-of-range artifact). Second, a low-frequency brightness change (low-frequency halo). It gives an unrealistic aspect to the result, especially when the contrast is reversed, *e.g.* a decreasing brightness from top to bottom becomes in the fused image an increasing brightness. A good illustration of it is given in [29]. Both artifacts are described in this section; illustration are shown in Figure 1, 2 and 3. Our goal in this section is to resolve these artifacts to get closer to this desirable result. We do not intend to fundamentally change EF since its results are already considered as the best available [26]. Our goal in the next Section 5 is to show that we can build on this improved technique a simple method to improve the contrast of single images.

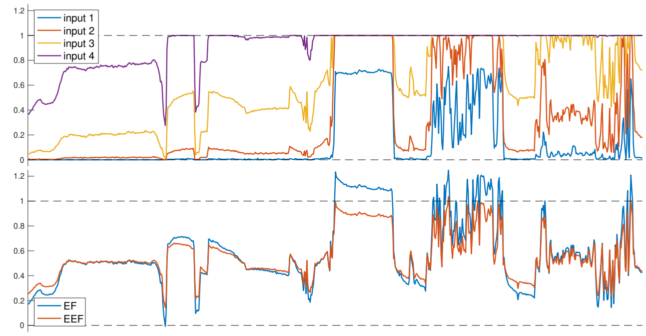

Figure 2. Out-of-range artifact. The sections displayed here are horizontal lines passing through the top left window of the house images. Top: input sequence (see Figure 1). Bottom: fused result with EF (blue) and EEF (red). While all input images are in the correct range, the EF result has a larger dynamic: this is the out-of-range artifact. The EEF method does not suffer from this artifact.

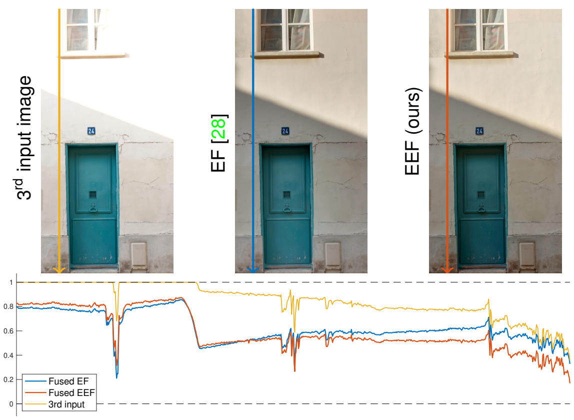

Figure 3. Low-frequency halo. The EF result has a vertical low-frequency halo. One of the input is displayed on the left for reference. In the EF result, the top of the shadow is darker than the bottom. That is, the brightness change is reverted with respect to the input images, where the top of the shadow is the brightest. This can be verified in the bottom plot where we drew lines from the images above. Our method (EEF) does not create this effect.

Out-of-range artifact While the sum of the weights is guaranteed for every pixel to be equal to 1, this does not imply that the reconstructed image belongs to the initial interval. This out-of-range artifact was described (yet not explained) by the Exposure Fusion's authors themselves [29]. Avoiding clipping in the dark or bright parts is possible by applying an affine rescaling of the image's dynamic to fit it to the standard interval $[0, 1]$. Yet the final compressed image may end up with a local contrast that is smaller than the contrast of the input images.

We illustrate this clipping/contrast dilemma in Figure 1, where we compare the standard (clipped) result with the normalized one. Numerous values of the standard output are lost, in the bright parts notably, *e.g.* the top left window pane. Normalizing the result so that it fits $[0, 1]$ gives an output with lower contrast than the input images. Figure 2 shows sections of the input images and a fused result.

This artifact happens because the method preserves the highly contrasted edges. But fitting every good part of the input images in the dynamic forces to reduce their amplitude [13]. Since EF only computes averages of Laplacian coefficients, the blending selects the edges with high amplitude, thus leading to an excessive image range.

Our proposed solution simply reduces the dynamic range of the input images. We shall see that this can be done in a way that preserves relevant information of each input image, and allows to reduce the final image edges' amplitude.

Extended exposure fusion Our correction of the out-of-range artifact in EF simply consists in separating each input image in a series of images with reduced dynamic range. To this aim, we use the function g defined in Equation (8). The only parameter is the restrained range width $\beta \in (0, 1]$. The number of generated images per input one is $M = \lceil 1/\beta \rceil$, where $\lceil \cdot \rceil$ is the ceiling function. This ensures that the full initial range is reproduced in the generated images.

First, let us call $\rho(k)$ the center of the reduced range, where $k \in \{-N^*, \dots, N\}$ is the index of the generated image in the series. For now we define $N = M - 1$ and $N^* = 0$, but these values will have different values in Section 5. The range is then limited to $[\rho(k) - \frac{\beta}{2}, \rho(k) + \frac{\beta}{2}]$, with

$$\rho(k) = 1 - \beta/2 - (k + N^*)(1 - \beta)/(N + N^*). \quad (7)$$

Rather than brutally clipping the values outside this restrained range, we use a function that progressively reduces the contrast until it becomes zero. We call g the function designed for this purpose:

$$g(t, k) = \begin{cases} t & \text{if } |t - \rho(k)| \leq \frac{\beta}{2} \\ \frac{t-\rho(k)}{|t-\rho(k)|}\left(a - \frac{\lambda^2}{|t-\rho(k)|-b}\right) + \rho(k) & \text{otherwise,} \end{cases} \quad (8)$$

where t denotes an intensity, $a = \frac{\beta}{2} + \lambda$ and $b = \frac{\beta}{2} - \lambda$, with λ is a parameter controling the speed of the decay in the function outside the valid part of the range. We keep it fixed to $\lambda = 0.125$. The decay behaves like $1/x^2$. These remapping functions are displayed in Figure 5. The particular shape of g is not important; any function with a sufficiently fast decay and a smooth transition to outside the valid range can be used. Smooth transitions avoid the creation of false edges that could be transmitted to the final image.

After fusion, the fused image's intensities need to be rescaled to the unit range. Depending on β and the content of the image, this typically amounts to stretch the intensity range (as opposed to EF). This final step over-stretches the colors to $[0, 1]$ by allowing 1% of clipping in both sides of the histogram. The method is described in Algorithm 2.

Algorithm 2: Extended Exposure Fusion (EEF)

```
input  : sequence û (with N images), reduced range β
output : z: fused image
K ← ⌈1/β⌉
for n ∈ {0, ..., N−1} do
    for k ∈ {0, ..., K} do
        v̂_{k+n×(K−1)} ← g(û_n, k)        // Equation (8)
z ← ExposureFusion(v̂)                    // Algorithm 1
```

Results and contrast enhancement We compare in Figure 4 the results obtained with EF (first column) and EEF, for different values of β. The remapping functions, used to distribute the initial intensities between the generated images, are displayed on the right. We observe that all EEF results with $\beta = 0.5$ are already more contrasted than the EF output. Hence, this value is sufficient to correct the artifact, as can be easily verified in Figure 2, where we showed an horizontal line in the "house", that goes through the top left pane of the windows, taken in the EF result (in blue) and the EEF result (in red, with $\beta = 0.48$). In this figure, the out-of-range artifact of the original method is obvious, and so is its correction with the proposed method.

The results with smaller β are even more contrasted locally. Thus, our correction can also be used to increase the local contrast with respect to the input images. Several methods have been proposed to produce more contrasted images with EF, *e.g.* by enhancing the local contrast in the sequence before the fusion [45] or after [22, 49]. With EEF, this can simply be obtained by setting β to a smaller value, and applying a final stretching step to the fused image.

Comparisons In addition to the comparison with Mertens *et al.* [28], we present in Figure 4 fusion results with the 2007 and 2009 methods of Raman *et al.* [35, 36] (MLSC and BBEF respectively), with Kou *et al.* [17] (ICME), Ma *et al.* [27] (PWMEF), and Li *et al.* [22] (DEEF). The MLSC and BBEF results are slightly darker and less contrasted locally than other methods, *e.g.* in the house image. The results with the MLSC method are a little blurry due to the regularization term on the fused output. The ICME method produces halos, that can easily be spotted in the sky of the grand canal image. The PWMEF method produces surprisingly bright results. It is visually pleasing, but one can notice out-of-range artifacts in the house image. Furthermore,

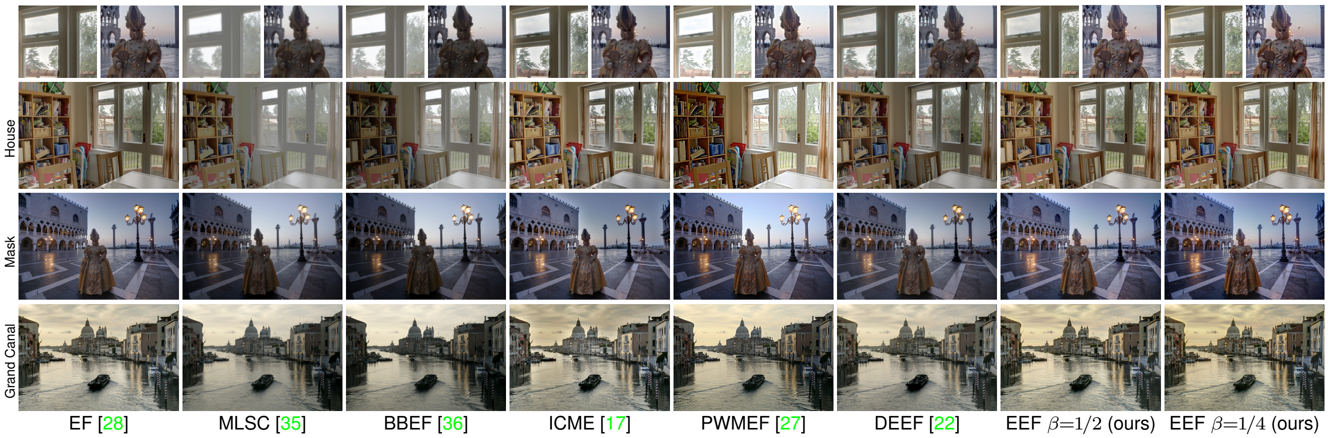

Figure 4. Comparison between the original EF, other fusion methods, and the proposed extended exposure fusion. All images were normalized with 0.1% of clipping in the white and 0.9% in the black. For EF, this results in a compression of the contrast, due to the out-of-range artifact, whereas for EEF this results in a stretching of the dynamic (the smaller β, the stronger the stretching).

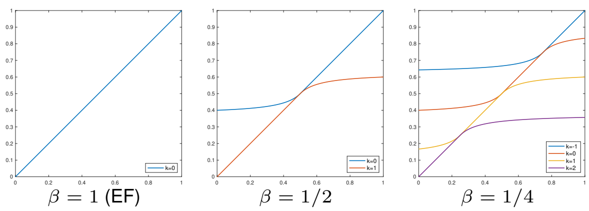

Figure 5. Remapping functions in EEF. $\beta = 1$ amounts to EF.

the bookshelf looks exaggeratedly saturated. Finally, the DEEF results are dark and lack local contrast. In fact, they look very similar to those of EF. EEF, however, is able to deliver fused images without halo and without out-of-range artifacts, and with an adequate amount of local contrast.

Elimation of the low-frequency halo The correction detailed above also resolves the second and last artifact known in Exposure Fusion: the low-frequency halo described in [29] and explained in [13] (see illustration in Figure 3). It appears near strong intensity differences in the input sequence. It corresponds to residual seams, but at the coarsest resolution of the pyramid. Its solution consists in using deeper pyramids. Yet, this trick could not be used in practice because it increased the out-of-range artifact [13]. With our correction, this is no longer an issue.

In Figure 6 we present an example of low-frequency halo. We created an image made of plateaus and a small noise to simulate texture. Two other images were generated using affine rescaling of the first one (and keeping the part in $[0, 1]$ only). Lines of the input sequence are displayed in the top plot, and fused results with EF and EEF and different pyramid depths in the center and bottom plots. The low-frequency halos of the EF method can be easily spotted on the left and right parts of the center blue line, where the standard depth is used. In EF the depth is computed as $\lfloor \log \min\{H, W\} / \log 2 \rfloor$, where H and W are the height and width of the image. Since these values are rarely powers of two, the residual image is generally a few pixels wide.

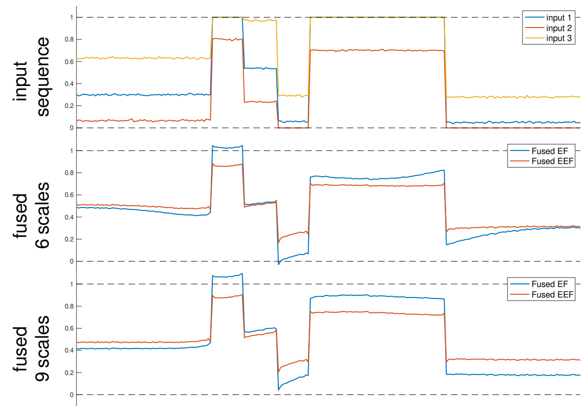

Figure 6. Fusion of a sequence (top) with the standard depth (center) or a larger depth (bottom), for EF (blue) and our EEF (red). Low-frequency halos are visible in the standard EF result. These halos can be removed using deeper pyramids, yet this also increases the out-of-range artifact. Our result, however, displays no halo nor out-of-range artifacts, independently of the depth.

The maximal depth, however, is obtained by continuing to downsample until the residual's size is one. In our experiment the Mertens's *et al*. depth is 6 and the maximal one is 9. These halos can be removed by using deeper pyramids [13, 29]. But the out-of-range artifact is amplified by this modification [13], as can be seen on the bottom plot. With EEF however, no low-frequency nor out-of-range artifacts are visible in the results.

In summary, with a simple modification of the algorithm, we have resolved the two identified shortcomings of the initial method: the out-of-range artifact, and the low-frequency halo, and provide a simple way to obtain fused images with improved local contrast.

Application to two large-exposure-ratio images

Recently Yang *et al*. presented a interesting variant of exposure fusion, the fusion of only two images with a large-

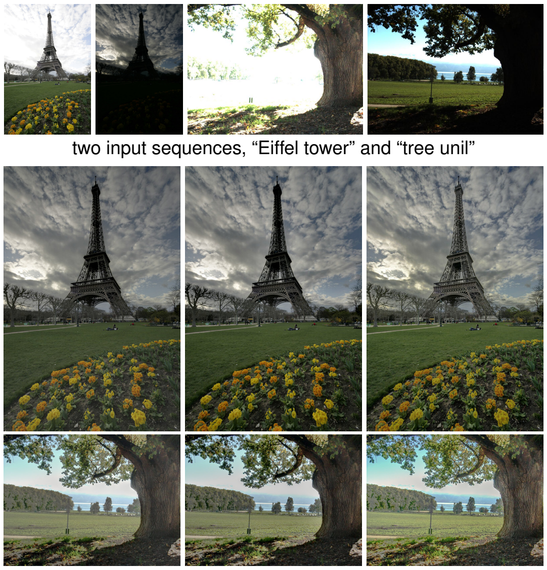

two input sequences, “Eiffel tower” and “tree unil”

EF [28] TIMEF [47] EEF (ours)

Figure 7. Fusion of a sequence of two images only, with a very large exposure time ratio. EEF handles this case without difficulty, and produces more contrasted results than TIMEF [47] that was specifically designed for this case. Our method is more general and produce visually better results.

exposure-ratio [47]. This makes sense for example when both images are captured at the same time with two sensors, or when using a video camera that alternatively captures long and short exposures. We show in Figure 7 that EEF can be applied in this configuration too, with a trivial modification: each simulated image is centered in 0.5. This avoids rejecting most regions, which are extremely bright or dark. Compared to the algorithm proposed in [47], our method achieves natural-looking images without changes in the relative brightness, and much more local contrast.

Limitations In areas where no input images contain information (*e.g.* all images are saturated), the EEF method will produce flat regions. See for example the saturated sky in the “tree unil” images in Figure 7. This limitation is shared with all multi-image fusion methods.

5. Simulated Exposure Fusion (SEF)

We present in this section our model for the simulation of a set of images from a single input. The simulated sequence is then fused with a refined extended exposure fusion. The pseudo-code of our method is given in Algorithm 3. Our aim is to enhance the contrast of a single image by simulating an increase of exposure time where needed. The desirable image is well-exposed in all regions.

Artificial sequence generation A bracketed exposure sequence is simulated from a single image using a simple set of contrast factors. Let us call $\alpha \geq 1$ the maximal contrast amplification factor. It is set by the user and controls the amount of enhancement. We shall define two remapping functions f^* and f to simulate darker and brighter images, respectively. We call them under- and over-exposed in the following developments, by a slight abuse of terminology. Under-exposed images are denoted with a negative index k. We call $N^* \in \mathbb{N}$ and $N \in \mathbb{N}^*$ the number of under- and over-exposed images respectively; their value is automatically computed. We also define $N_{\max} = \max(N, N^*)$. The bracketed sequence is then simulated using:

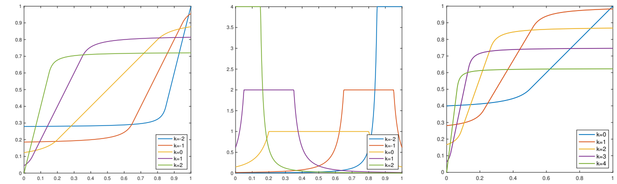

(a) remapping functions with $N^* = N$ (b) derivatives of (a) (c) remapping functions with $N^* < N$

Figure 8. Remapping functions and their derivatives. The restrained dynamic range location is adapted to the different generated images, so a to preserve the most pertinent information. The number of under- (N^*) and over-exposed (N) images is automatically computed from the histogram.

$$\begin{cases} f^*(t,k) = \alpha^{|k|/N_{\max}}(t-1)+1 & \text{if } k<0, \\ f(t,k) = \alpha^{k/N_{\max}}t & \text{otherwise.} \end{cases} \tag{9}$$

The function f^* increases the contrast like f, but also shifts the intensities toward the bottom to prevent clipping of the bright values (we recall that $t \in [0,1]$). The unmodified input image is included in the generated sequence and denoted with index $k = 0$. The simulated sequence has then $M = N + N^* + 1$ images, numbered from $-N^*$ to N. Images whose indices have the same absolute value have the same enhancement factor. Because the factors are equal or superior to 1, the fused image is guaranteed to gain contrast. Higher contrast factors are applied to the far left and right side of the histogram, where more enhancement is needed.

In the raw case, increasing the exposure time amounts to applying a factor to the image. Although in the JPEG case non-linear corrections (gamma and s-shaped curve) have been applied, this equivalence is still approximately true locally in the range. This justifies the application of simple contrast factors in Equation (9), since they are used in limited portion of the range only thanks to Equation (10).

Adopting the restrained dynamic range strategy of EEF, the enhanced images are generated using the composition of the two functions, *i.e.* using

$$\begin{cases} (g \circ f^*)(t,k) & \text{if } k<0 \\ (g \circ f)(t,k) & \text{otherwise.} \end{cases} \tag{10}$$

These functions are displayed in Figure 8 (a) and (c).

To show the effect of the dynamic reduction on the fusion result, we did the following experiment. We created an image by repeating the same 1D signal along the lines. This is the blue one in the bottom plots (a) and (b) in Figure 9.

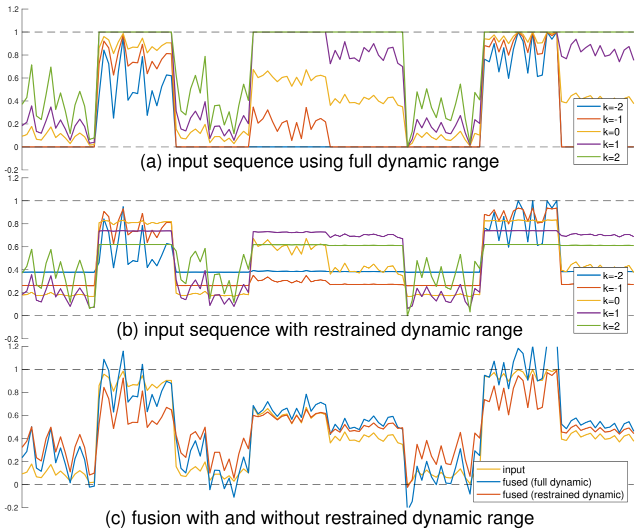

Figure 9. Comparison of the fused results without (a) and with (b) the restrained dynamic. In (b), we used $\beta = 0.5$. Comparing the left and right results one can check that the out-of-range values have been almost completely removed.

From this image we generated two bracketed exposure sequences, without and with dynamic reduction, *i.e.*, using Equation (9) and (10), respectively. They are displayed on the top left and right columns. We then fused them. The line from the resulting image is superimposed, in red, in the bottom plots (a) and (b). The out-of-range artifact of the method is clearly visible in the fused result (a). However, it is absent from the fused result (b). This illustrates the effectiveness of our proposition. Observe that the main edges are generally not reduced in (a), whereas they are strongly reduced in (b). Hence in SEF too the restrained dynamic strategy solves EF's inherent out-of-range artifact.

An image rarely needs the same enhancement in its bright and dark areas. In fact, contrast enhancement is generally needed in the dark parts only, *e.g.* in the case of backlit images. Enhancement in the bright parts can still be beneficial to thwart the effect of gamma-correction and S-shaped curves that compress the contrast in this region of the intensity range. The number of darkened or brightened images N and N^* is thus determined from the input image's histogram, using the median. N^* is computed using M as: $N^* = \lfloor (M-1)\text{median}\{u\} \rfloor$, where $\lfloor r \rfloor$ denotes the integer part of r. The number of over-exposed images is then $N = M - N^* - 1$. The value M is automatically computed from the user-set parameters α and β by ensuring that all parts of the dynamic are enhanced at least once in the simulated sequence.

HSV colorspace Once the bracketed exposure sequence is simulated, it is fused with a refined EF. Notably, instead of modifying the color image, we convert the input to the HSV color space [42] and enhance the luminance only, *i.e.* V (value), while preserving the H (hue) and S (saturation) channels. This color space has the advantage of preserving the dynamic range when restoring the color, whereas other color spaces tend to generate colors outside the RGB color cube. Besides, this accelerates the algorithm. We thus do not use the saturation measure proposed by Mertens *et al.*

Contrast metric Knowing the remapping functions, we can easily and efficiently compute the contrast metric. It is directly given by the derivative of the remapping function:

$$\begin{cases} \alpha^{k/N_{\max}}(g' \circ f)(t,k) & \text{if } k < 0 \\ \alpha^{|k|/N_{\max}}(g' \circ f^*)(t,k) & \text{otherwise,} \end{cases} \tag{11}$$

where g' is the derivative of g w.r.t. t:

$$g'(t,k) = \begin{cases} 1 & \text{if } |t-\rho(k)| \leq \frac{\beta}{2} \\ \frac{\lambda^2}{(|t-\rho(k)|-b)^2} & \text{otherwise.} \end{cases} \tag{12}$$

Figure 8 (b) shows an example of derivatives. The well-exposedness measure is kept unchanged.

Algorithm 3: Simulated Exposure Fusion (SEF)

input : image u and parameters α, β
output : enhanced image v
$\ell \leftarrow$ V channel of input in HSV color space
while *all dynamic not covered* **do**
  $M \leftarrow M + 1$ `// Initialized at 1`
  $N^* \leftarrow \lfloor (M-1)\text{med}\{\ell\} \rfloor$; $N \leftarrow M - N^* - 1$
for $k \in \{-N^*, \dots, N\}$ **do** `// Generate a sequence`
  $\widehat{\ell}_k \leftarrow$ GENERATIONMODEL(ℓ) `// Equ. (10)`
  $w_{c,k} \leftarrow$ CONTRASTWEIGTHS(ℓ) `// Equ. (11)`
$w_e \leftarrow$ compute well-exposedness `// Equation (3)`
$\widehat{w} \leftarrow$ compute final weight maps `// Eq. (4) and (5)`
$\ell' \leftarrow$ multiscale fusion of $\widehat{\ell}$ using $\widehat{w}$ `// Equation (6)`
$v_{\text{RGB}} \leftarrow$ replace V of u_{HSV} by ℓ', convert back to RGB

6. Results

We compared in Figure 10 our results with other retinex and histogram-based methods: Automatic Color Enhancement (ACE) [5, 6], Multiscale Retinex (MSR) [18, 33], Histogram Equalization (HE), Contrast Limited Adaptive HE (CLAHE) [34], Detail-Enhanced Exposure Fusion (DEEF) [22], and Fast Local Laplacian Filter (LLF) [1, 31].

The results of SEF are comparatively more definite. In the "columns" image for example, the SEF result (first line) is the only one that correctly enhances the bushes and the house wall between the columns. Furthermore, SEF did not create halos, contrarily to CLAHE (and also, to a lesser extent MSR and LLF). It also did not lose contrast anywhere in the images, contrarily to HE, ACE, DEEF and MSR. Unlike CLAHE, SEF preserves the global contrast: for example, the sky of the "column" image is not darkened. Also, SEF does not over-enhance some regions that do not need it, contrarily to LLF in the "escanonge" image for example, that also enhances the already correctly exposed parts of the image, resulting in too much local contrast on the sky.

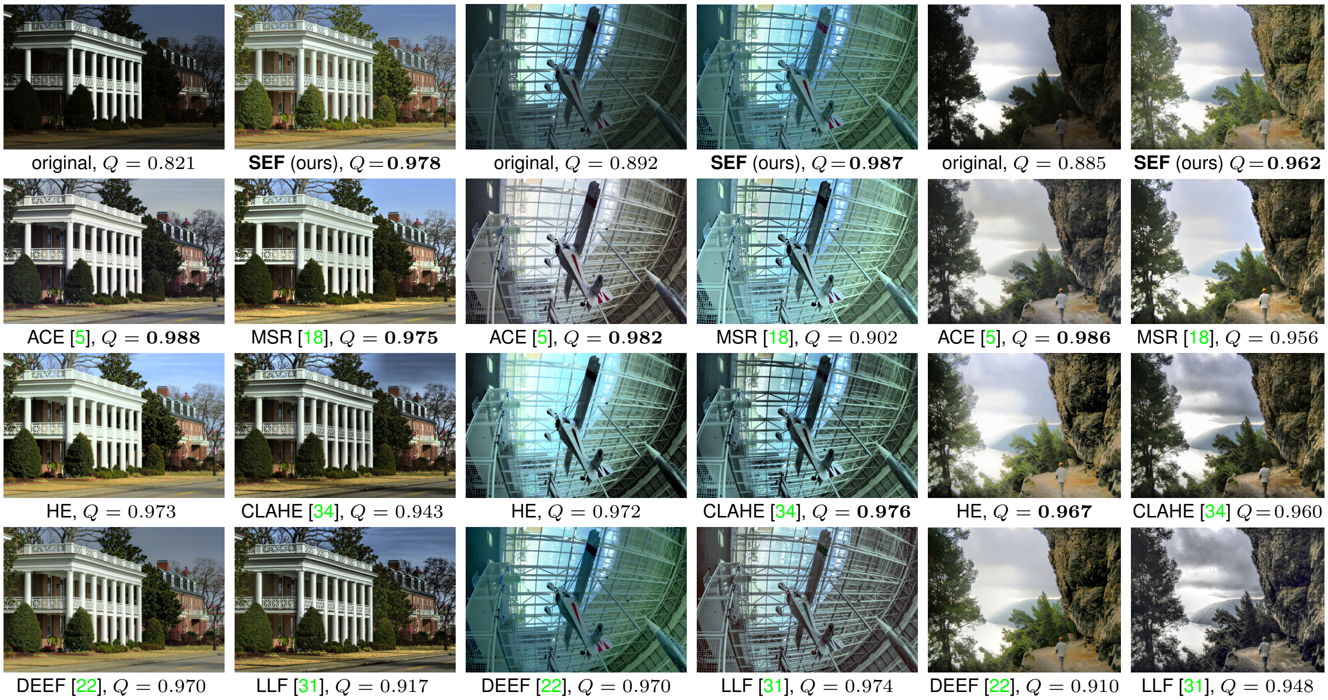

Figure 10. Comparison of SEF with other retinex methods on three images. SEF, ACE and LLF have an enhancement parameter (called α for SEF); it is set to 8. Similarly to SEF, the algorithms HE, CLAHE, MSR and LLF were applied to the luminance channel only. For SEF, we used $\beta = 0.5$ which gives $N = 4$ and $N^* = 0$: only five images were fused. For LLF, we used the default parameters $\sigma = 0.1$, $N = 10$, which means that 10 images were fused. For ACE and MSR we used the default parameters proposed in [6, 33]. For DEEF we used the authors' implementation default parameters. For HE and CLAHE we used the Matlab implementation default parameters.

DEEF does not enhance bright regions. Besides, it produces images with values standing largely outside $[0, 1]$ (the out-of-range artifact), thus requiring a final range compression. This means that bright regions systematically lose contrast. This phenomenon is apparent in the sky of the "columns" and "escanonge" images. Moreover, the dark parts in DEEF are not as well enhanced as with SEF.

We reported in Figure 10 the tone-mapping quality index (TMQI) [48] score of the enhanced images. In each case our method is among the three best ones.

In short, the proposed method outperforms state-of-the-art retinex algorithms such as MSR, ACE, and CLAHE, because it preserves and enhances contrast everywhere without creating halos. As for DEEF [22], the method suffers from the out-of-range artifact, which produces images with very poor local contrast. Besides, DEEF is more complex and costly since it requires edge-aware smoothing in addition to the Laplacian pyramid blending. Concerning edge-aware local contrast enhancement, the state-of-the-art LLF method is able to correctly enhance contrast everywhere without creating artifacts (slight halos are sometimes visible). Compared to LLF, our algorithm produces more natural images, because it corrects the exposition locally rather than enhancing local contrast with the same factor everywhere. It therefore better adapts to the input image. This also saves computations. In Figure 10, SEF only required three Laplacian pyramids, while LLF used ten.

Limitations The single image contrast enhancement technique SEF increases noise as much as contrast.

7. Conclusion

We have shown that the out-of-range artifact and the low-frequency halo in Exposure Fusion originated in the lack of a mechanism to reduce strong edges. Limiting the range of the input resolves these drawbacks and produces results free of halos and clipped values, and with improved local contrast everywhere. The resulting method, which we called extended exposure fusion (EEF), produces better results than concurrent approaches. Applied to a single image the method simulates first under- and over-exposed images, then applies EEF to them. This yields a state of the art LDR contrast enhancement method that we called simulated exposure fusion (SEF). The proposed improvements are easily applicable to all methods involving exposure fusion.

Acknowledgements Work partly financed by Office of Naval research grant N00014-17-1-2552, DGA Astrid project "filmer la Terre" number ANR-17-ASTR-0013-01. Charles Hessel's PhD was supported by a CIFRE scholarship of the French Ministry for Higher Studies, Research and Innovation. He wishes to thank his DxO collaborators for many valuable conversations and advice.

References

[1] M. Aubry, S. Paris, S. W. Hasinoff, J. Kautz, and F. Durand. Fast and robust pyramid-based image processing. Technical report, Computer Science and Artificial Intelligence Laboratory, Nov. 2011. 2, 7

[2] L. Bogoni. Extending dynamic range of monochrome and color images through fusion. In *Proceedings 15th International Conference on Pattern Recognition. ICPR-2000*, volume 3, pages 7–12 vol.3, 2000. 1, 2

[3] P. Burt and E. Adelson. The laplacian pyramid as a compact image code. *IEEE Transactions on communications*, 31(4):532–540, 1983. 3

[4] P. J. Burt and R. J. Kolczynski. Enhanced image capture through fusion. In *Computer Vision, 1993. Proceedings., Fourth International Conference on*, pages 173–182. IEEE, 1993. 1, 2

[5] C. Gatta, A. Rizzi, and D. Marini. Ace: An automatic color equalization algorithm. In *Conference on Colour in Graphics, Imaging, and Vision*, volume 2002, pages 316–320. Society for Imaging Science and Technology, 2002. 2, 7, 8

[6] P. Getreuer. Automatic color enhancement (ace) and its fast implementation. *Image Processing On Line*, 2:266–277, 2012. 2, 7, 8

[7] J. G. Gomila Salas and J. L. Lisani. Local Color Correction. *Image Processing On Line*, 1:260–280, 2011. 2

[8] A. A. Goshtasby. Fusion of multi-exposure images. *Image and Vision Computing*, 23(6):611 – 618, 2005. 1

[9] B. Gu, W. Li, J. Wong, M. Zhu, and M. Wang. Gradient field multi-exposure images fusion for high dynamic range image visualization. *Journal of Visual Communication and Image Representation*, 23(4):604–610, may 2012. 2

[10] D. Hafner and J. Weickert. Variational Image Fusion with Optimal Local Contrast. *Computer Graphics Forum*, 35(1):100–112, 2016. 1

[11] S. W. Hasinoff, D. Sharlet, R. Geiss, A. Adams, J. T. Barron, F. Kainz, J. Chen, and M. Levoy. Burst photography for high dynamic range and low-light imaging on mobile cameras. *ACM Transactions on Graphics*, 35(6):1–12, nov 2016. 2

[12] J. Herwig and J. Pauli. An information-theoretic approach to multi-exposure fusion via statistical filtering using local entropy. In *Proc. International Conference on Signal Processing, Pattern Recognition and Applications*, pages 50–57, 2010. 1, 2

[13] C. Hessel. An Implementation of the Exposure Fusion Algorithm. *Image Processing On Line*, 8:369–387, 2018. 4, 5

[14] W. ho Cho and K.-S. Hong. Extending dynamic range of two color images under different exposures. In *Proceedings of the 17th International Conference on Pattern Recognition, 2004. ICPR 2004.*, volume 4, pages 853–856 Vol.4, Aug 2004. 2

[15] D. J. Jobson, Z.-u. Rahman, and G. A. Woodell. A multiscale retinex for bridging the gap between color images and the human observation of scenes. *IEEE Transactions on Image processing*, 6(7):965–976, 1997. 2

[16] K. Kotwal and S. Chaudhuri. An optimization-based approach to fusion of multi-exposure, low dynamic range images. In *14th International Conference on Information Fusion*, pages 1–7, July 2011. 1

[17] F. Kou, Z. Li, C. Wen, and W. Chen. Multi-scale exposure fusion via gradient domain guided image filtering. In *Proceedings - IEEE International Conference on Multimedia and Expo*, number July, pages 1105–1110. IEEE, jul 2017. 2, 4, 5

[18] E. H. Land. The retinex theory of color vision. *Scientific American*, 237(6):108–129, 1977. 2, 7, 8

[19] C.-H. Lee, L.-H. Chen, and W.-K. Wang. Image contrast enhancement using classified virtual exposure image fusion. *IEEE Transactions on Consumer Electronics*, 58(4):1253–1261, nov 2012. 2

[20] S. Li and X. Kang. Fast multi-exposure image fusion with median filter and recursive filter. *IEEE Transactions on Consumer Electronics*, 58(2):626–632, may 2012. 2

[21] S. Li, X. Kang, and J. Hu. Image Fusion With Guided Filtering. *IEEE Transactions on Image Processing*, 22(7):2864–2875, jul 2013. 1, 2

[22] Z. Li, Z. Wei, C. Wen, and J. Zheng. Detail-Enhanced Multi-Scale Exposure Fusion. *IEEE Transactions on Image Processing*, 26(3):1243–1252, mar 2017. 2, 4, 5, 7, 8

[23] Z. Li, J. Zheng, Z. Zhu, W. Yao, and S. Wu. Weighted guided image filtering. *IEEE Transactions on Image Processing*, 24(1):120–129, 2015. 2

[24] J. L. Lisani. Adaptive Local Image Enhancement Based on Logarithmic Mappings. In *2018 25th IEEE International Conference on Image Processing (ICIP)*, volume 255, pages 1747–1751. IEEE, oct 2018. 2

[25] K. Ma, Z. Duanmu, H. Yeganeh, and Z. Wang. Multi-Exposure Image Fusion by Optimizing A Structural Similarity Index. *IEEE Transactions on Computational Imaging*, 4(1):60–72, mar 2018. 1

[26] K. Ma, Kai Zeng, and Zhou Wang. Perceptual Quality Assessment for Multi-Exposure Image Fusion. *IEEE Transactions on Image Processing*, 24(11):3345–3356, nov 2015. 2, 3

[27] K. Ma and Z. Wang. Multi-exposure image fusion: A patch-wise approach. In *2015 IEEE International Conference on Image Processing (ICIP)*, pages 1717–1721. IEEE, sep 2015. 1, 4, 5

[28] T. Mertens, J. Kautz, and F. Van Reeth. Exposure fusion. In *Computer Graphics and Applications, 2007. PG '07. 15th Pacific Conference on*, pages 382–390, Oct 2007. 1, 2, 3, 4, 5, 6

[29] T. Mertens, J. Kautz, and F. Van Reeth. Exposure fusion: A simple and practical alternative to high dynamic range photography. *Computer Graphics Forum*, 28(1):161–171, 2009. 1, 2, 3, 4, 5

[30] J. M. Ogden, E. H. Adelson, J. R. Bergen, and P. J. Burt. Pyramid-based computer graphics. *RCA Engineer*, 30(5):4–15, 1985. 3

[31] S. Paris, S. W. Hasinoff, and J. Kautz. Local laplacian filters: edge-aware image processing with a laplacian pyramid. *ACM Trans. Graph.*, 30(4):68, 2011. 2, 7, 8

[32] S. Paris, S. W. Hasinoff, and J. Kautz. Local laplacian filters: edge-aware image processing with a laplacian pyramid. *Communications of the ACM*, 58(3):81–91, 2015. 2

[33] A. B. Petro, C. Sbert, and J.-M. Morel. Multiscale retinex. *Image Processing On Line*, pages 71–88, 2014. 2, 7, 8

[34] S. M. Pizer, E. P. Amburn, J. D. Austin, R. Cromartie, A. Geselowitz, T. Greer, B. ter Haar Romeny, J. B. Zimmerman, and K. Zuiderveld. Adaptive histogram equalization and its variations. *Computer Vision, Graphics, and Image Processing*, 39(3):355–368, sep 1987. 7, 8

[35] S. Raman and S. Chaudhuri. A matte-less, variational approach to automatic scene compositing. In *2007 IEEE 11th International Conference on Computer Vision*, pages 1–6, Oct 2007. 1, 4, 5

[36] S. Raman and S. Chaudhuri. Bilateral filter based compositing for variable exposure photography. In *Eurographics (Short Papers)*, pages 1–4, 2009. 1, 2, 4, 5

[37] E. Reinhard, W. Heidrich, P. Debevec, S. Pattanaik, G. Ward, and K. Myszkowski. *High dynamic range imaging: acquisition, display, and image-based lighting*. Morgan Kaufmann, 2010. 1

[38] A. Rizzi, C. Gatta, and D. Marini. A new algorithm for unsupervised global and local color correction. *Pattern Recognition Letters*, 24(11):1663–1677, 2003. 2

[39] S. Rui, I. Cheng, S. Jianbo, and A. Basu. Generalized Random Walks for Fusion of Multi-Exposure Images. *IEEE Transactions on Image Processing*, 20(12):3634–3646, dec 2011. 1, 2

[40] R. Shen, I. Cheng, and A. Basu. Qoe-based multi-exposure fusion in hierarchical multivariate gaussian crf. *IEEE Transactions on Image Processing*, 22(6):2469–2478, June 2013. 2

[41] H. Singh, V. Kumar, and S. Bhooshan. Weighted least squares based detail enhanced exposure fusion. *ISRN Signal Processing*, 2014, 2014. 1, 2

[42] A. R. Smith. Color gamut transform pairs. *ACM SIGGRAPH Computer Graphics*, 12(3):12–19, aug 1978. 7

[43] M. Song, D. Tao, C. Chen, J. Bu, J. Luo, and C. Zhang. Probabilistic exposure fusion. *IEEE Transactions on Image Processing*, 21(1):341–357, Jan 2012. 2

[44] C. Wang and C. Tu. Different Levels of Detail Display for Exposure Fusion Using Local Laplacian Filtering. *Journal of Computers*, 9(9):2181–2185, 2014. 2

[45] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004. 1, 4

[46] Wei Zhang and Wai-Kuen Cham. Gradient-Directed Multiexposure Composition. *IEEE Transactions on Image Processing*, 21(4):2318–2323, apr 2012. 2

[47] Y. Yang, W. Cao, S. Wu, and Z. Li. Multi-Scale Fusion of Two Large-Exposure-Ratio Images. *IEEE Signal Processing Letters*, 25(12):1885–1889, dec 2018. 6

[48] H. Yeganeh and Z. Wang. Objective quality assessment of tone-mapped images. *IEEE Transactions on Image Processing*, 22(2):657–667, 2013. 8

[49] Zheng Guo Li, Jing Hong Zheng, and S. Rahardja. Detail-Enhanced Exposure Fusion. *IEEE Transactions on Image Processing*, 21(11):4672–4676, nov 2012. 2, 4