

Published in Image Processing On Line on 2019–12–29. Submitted on 2017–09–17, accepted on 2019–12–28. ISSN 2105–1232 © 2019 IPOL & the authors CC–BY–NC–SA This article is available online with supplementary materials, software, datasets and online demo at https://doi.org/10.5201/ipol.2019.279

Simulated Exposure Fusion

Charles Hessel

CMLA, ENS Cachan, CNRS, Université Paris-Saclay, 94235 Cachan, France

Abstract

Simulated Exposure Fusion (SEF) is a single-image contrast enhancement method. It is built upon a high dynamic range imaging technique called Exposure Fusion (EF), introduced in 2007 and widely used since then, which aims at fusing a bracketed exposure sequence into a high quality image. Simulated Exposure Fusion extends the initial method to the case where only one image is available, and delivers an image with enhanced contrast. We propose in this paper an implementation of this method, along with its precise description and analysis. Its results are compared to state-of-the-art enhancement algorithms and appear to be artifact-free, even in extreme enhancement conditions. Furthermore, they inherit from EF's celebrated natural aspect.

Source Code

The Matlab/Octave source code, the code documentation, and the online demo are accessible at the web page of this article¹. Usage instructions are included in the README.txt file of the archive.

Keywords: high dynamic range; image fusion; tone-mapping; Laplacian pyramids

1 Introduction

Exposure Fusion (EF) is a high dynamic range imaging technique to fuse a bracketed exposure sequence into a high quality image, introduced in 2007 by Mertens et al. [14, 15]. 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. A precise description of the method along with its main limitation, namely, an out-of-range artifact can be found in [7].

In this paper, following the lines of the method introduced in [8], we describe a single image contrast enhancement method derived from EF, called Simulated Exposure Fusion (SEF). Starting from a single image, SEF simulates a bracketed exposure sequence of images and fuses them by an adaptation of EF avoiding its out-of-range artifact.

¹https://doi.org/10.5201/ipol.2019.279

Our plan is as follows. We start by presenting in Section 2 the Simulated Exposure Fusion (SEF). Its two main steps are presented in two distinct sections: Section 2.1 for the sequence generation and Section 2.2 for the few modifications to the initial fusion method of Mertens et al. [14]. A study of the effect of the parameters of the method, as well as a comparison with state-of-the-art contrast enhancement methods, are carried out in Section 3.

2 Simulated Exposure Fusion

This section presents SEF [8], a method generating a set of images with simulated exposure from a single input. This model simply applies different multiplicative factors to the input image. The simulated sequence is then fused with a refined exposure fusion. The whole process is displayed in Figure 1. The pseudo-code of the method is given in Algorithm 2.



Figure 1: Simulated exposure fusion method (SEF).

2.1 Generation of an Artificial Bracketed Exposure Sequence

The simulated exposure fusion method uses remapping functions to generate an artificial bracketed exposure sequence. Arguably, the choice of these functions is the most important point of the method. As they are defined, they include two major ideas: the generation should always increase the contrast, and the dynamic should be *restrained*. In practice, this is implemented with two functions, f and g, that are composed, so that the simulated sequence is obtained using

$$\begin{cases} (g \circ f^*)(t,k), & \text{if } k < 0\\ (g \circ f)(t,k), & \text{otherwise,} \end{cases}$$
(1)

where t is an intensity and k the position of the image in the sequence; over-exposed images are denoted with a positive index and under-exposed images with a negative one. These functions depend on three parameters. Two are user-set and the third depends on the input image's histogram.

- The first parameter, α , controls the strength of the enhancement; it is linked to the generation model as it amounts to a factor on the exposure time of the input image. It intervenes in function f.
- The second parameter, β , controls the dynamic range compression, and thus intervenes in function g. It is the key that solves the out-of-range artifact of the initial EF method. It implicitly reduces the edges' amplitude, hence it can also be seen as a way to control (limit) the amplitude of the structures that are being enhanced.

• The last parameter, m, is automatically set as the median of the input image. On its value depend the numbers N and N^* of over- and under-exposed images in the simulated sequence. Hence m controls the brightening or darkening of the image in the simulated sequence. As a result, the enhanced image's histogram is always more centered than the histogram of the input.

We shall present each of them in more detail in the three next Sections 2.1.1, 2.1.2 and 2.1.3. Figure 2 displays the set of remapping functions obtained in function of the values of the two parameters α and β , and of the median m of the image. We shall follow in this paper the convention that the images' displayable range is [0, 1].



Figure 2: Remapping functions in function of the three parameters on which they depend. We denote $median\{u\}$ by m.

2.1.1 The Remapping Functions

Although the simulated images in [8] are obtained using a set of contrast factors, the authors call them under- and over-exposed. Indeed, using a simplistic model of the digital images formation (and notably omitting the non-linearities of the acquisition process), simulating images with a shorter or longer exposure amounts to multiply them by a constant.

Let us denote by u the input image, by \mathbf{x} a pixel, and by \hat{u} a simulated image. To artificially increase the exposure time, the authors of [8] use

$$\widehat{u}_k(\mathbf{x}) = \min\{\lambda^k u(\mathbf{x}), 1\}, \qquad (2)$$

where $\lambda \in \mathbb{R}^+$ is a contrast enhancement factor set to the power $k \in \mathbb{Z}$, which is also the index of the image in the sequence. For over-exposed images, the authors set $\lambda^k \geq 1$. The contrast-enhanced image is clipped to fit the standard [0, 1] dynamic range.

Using $\lambda^k \geq 1$ can only generate over-exposed images from the input, thus the dark parts will be enhanced but the bright ones will always be clipped. However, simulating images with a shorter exposure time, i.e. using $\lambda^k < 1$, is not desirable, because they would have lower contrast and still no information in the clipped pixels. The authors of [8] therefore suggest to generate images with improved contrast in the bright areas by shifting the intensities toward the bottom, so that bright parts are enhanced and the dark parts clipped (this will become clear in Equation (3) below). By an abuse of terminology, the authors of [8] call these images under-exposed.

Remapping functions Let us call f the function used to generate the $N \in \mathbb{N}^*$ over-exposed images, and f^* the function to generate the $N^* \in \mathbb{N}^*$ under-exposed images. The values N and N^* are automatically computed, as described in Section 2.1.3. They may be different, we thus also define $N_{\max} = \max(N, N^*)$. The value of λ is deduced from N_{\max} and $\alpha \ge 1$, a user-set parameter giving the maximal contrast amplification factor. It controls the global amount of enhancement. The maximal factor α is reached with the last images of the sequence, i.e. when $k = N_{\max}$, thus $\lambda = \alpha^{1/N_{\max}}$.

The clipping is removed since this will be the purpose of function g, introduced latter. We recall that the under-exposed images are denoted with a negative index k, and the over-exposed with a positive index. 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. The reason is that images whose indices have the same absolute value have the same contrast enhancement factor. Denoting t an intensity, the sequence $\{\hat{u}_k : k = -N^*, \ldots, N\}$ is simulated with functions

$$\begin{cases} f^*(t,k) = \alpha^{|k|/N_{\max}}(t-1) + 1, & \text{if } k < 0, \\ f(t,k) = \alpha^{k/N_{\max}}t, & \text{otherwise.} \end{cases}$$
(3)

Because the multiplicative factors are equal or superior to 1, the fused image is guaranteed not to lose contrast. The first parameter of the remapping functions, α , then allows to controls the global amount of enhancement. It is equivalent to an increase of the exposure time.

2.1.2 Restraining the Dynamic of the Simulated Images

The simulated sequence is then fused with an exposure fusion technique. In [8] is proposed a modification of the fusion method of Mertens et al. [15] which fixes its out-of-range artifacts. The idea is simply to reduce the dynamic range of the input images in the generated sequence, so that the increased dynamic range of the fused image still fits the standard [0,1] range. This correction was proposed in [8] for the extended exposure fusion (EEF) algorithm and then adapted to the simulated exposure fusion that we describe here. This operation, which implicitly amounts to reduce the edges' amplitude, is called "restraining the range" by the authors, so as to avoid a confusion with the unwanted clipping of out-of-range values required at the end of the original EF. This is done in a way that preserves relevant information of each input image.

Restrained dynamic range In order to reduce the range of the generated images while keeping enough relevant information, the remapping functions are adapted to the generated images: the method preserves the initially dark regions in the images simulating an increase of exposure time, and preserves the bright regions in the simulated under-exposed images.

Let us call $\beta \in (0, 1]$ the restrained dynamic range. The method aims at preserving the values in $[0, \beta]$ for k = N and in $[1 - \beta, 1]$ for $k = -N^*$, and compressing the remaining part of the dynamic. Let us call $\rho(k)$ the center of the reduced range. The authorized range of any image k is then limited to $[\rho(k) - \frac{\beta}{2}, \rho(k) + \frac{\beta}{2}]$, where ρ is defined as

$$\rho(k) = 1 - \frac{\beta}{2} - \frac{(k+N^*)(1-\beta)}{N+N^*}.$$
(4)

The value of β is a parameter set by the user. Rather than brutally clipping the values outside the range $[\rho(k) - \frac{\beta}{2}, \rho(k) + \frac{\beta}{2}]$, the method opts for a "smooth clipping", using a function that will progressively reduce the contrast until it becomes zero. We call g the function designed for this purpose

$$g(t,k) = \begin{cases} t, & \text{if } |t - \rho(k)| \le \frac{\beta}{2} \\ \operatorname{sign}(t - \rho(k)) \left(a - \frac{\lambda^2}{|t - \rho(k)| - b} \right) + \rho(k), & \text{otherwise,} \end{cases}$$
(5)

where the values a and b are defined such that $a = \frac{\beta}{2} + \lambda$ and $b = \frac{\beta}{2} - \lambda$, where λ is a parameter that controls the speed of the decay in the function outside the "valid" part of the range. Its value is fixed to $\lambda = 0.125$. The decay behaves like $1/x^2$.

The enhanced images are hence generated using the composition of the two functions as described in Equation (1). Examples of these functions are displayed in Figure 2.

2.1.3 Automatic and Asymmetric Bracketed Exposure

The set of the remapping functions defined in Equation (1) relies on a third value: the median of the input image. This value helps defining a sequence adapted to the input's histogram, by means of the simulation of different numbers N and N^* of over- and under-exposed images, respectively. The median value is used in [8], because it gives a good estimation of the proportion of dark and bright pixels. The number of under-exposed images N^* is computed using M as

$$N^* = \lfloor (M-1) \operatorname{median}\{u\} \rfloor, \tag{6}$$

where $\lfloor r \rfloor$ denotes the integer part of r and M > 1. The number of over-exposed images is then $N = M - N^* - 1$.

Nonetheless, these formulas do not give the actual number of over- and under-exposed images, since M is unknown. Yet, following [8], the value of M can be deduced from the two user-set parameters α and β by ensuring that all parts of the dynamic are enhanced at least once in the simulated sequence. A small routine is hence proposed by the authors to solve this minimization problem that can be expressed as follows: find the smallest value M > 1 such that the following two conditions are satisfied for all k

$$\begin{cases} f^{\diamond}(\rho(k) + \frac{\beta}{2}, k) > f^{\diamond}(\rho(k-1) - \frac{\beta}{2}, k-1), & \forall k > -N^{*}, \\ f^{\diamond}(\rho(k) - \frac{\beta}{2}, k) < f^{\diamond}(\rho(k+1) + \frac{\beta}{2}, k+1), & \forall k < N, \end{cases}$$
(7)

where f^{\diamond} stands for f^* if k < 0 and f otherwise. This amounts to requiring a minimal overlap between all "valid" ranges of the remapping functions. In practice, this constraint can be computed for the images with index k = -1, 0 and +1 only, because they are the ones with the smallest overlap. The procedure is given in Algorithm 1.

2.2 Merging the Simulated Sequence

Once the sequence generated, the images are merged by exposure fusion. The authors in [8] however did three modifications to the initial method of Mertens et al. [14], which are:

- 1. working in gray, hence suppressing the "saturation" measure;
- 2. deducing the contrast measure from the remapping functions;
- 3. adding the possibility to use deeper pyramids.

Algorithm (1: Find	l number	of images	required in	the sequence	(SEQUENCELENGTH)
0				- 1		

input : Parameters α , β and $m = \text{median}\{u\}$ output: M, N^*, N 1 do $M \leftarrow M + 1$ $\mathbf{2}$ // initialized at 1 $N^* \leftarrow |m(M-1)|$ // Equation (6)3 $N \leftarrow (\tilde{M} - 1) - \tilde{N^*}$ 4 $N_{\max} \leftarrow \max\{N^*, N\}$ 5 $t_{\max,k=+1} = \alpha^{1/N_{\max}} \left(1 + \frac{\beta - 1}{M - 1} (N^* + 1)\right)$ $t_{\min,k=-1} = \alpha^{1/N_{\max}} \left(-\beta + \frac{\beta - 1}{M - 1} (N^* - 1)\right) + 1$ $t_{\max,k=0} = N^* \frac{\beta - 1}{M - 1} + 1$ $t_{\min,k=0} = N^* \frac{\beta - 1}{M - 1} + 1 - \beta$ // maximal valid value for k=+16 // minimal valid value for k=-17 // maximal valid value for k=08 // minimal valid value for k=09 10 while $t_{\max, k=+1} < t_{\min, k=0}$ or $t_{\max, k=0} < t_{\min, k=-1}$

2.2.1 Luminance Only

Contrast enhancement and local exposure correction in the SEF method are performed on the luminance, while preserving the color ratios. This is justified by a significant drop in computational cost and memory footprint. Furthermore, the authors argue that the result is unchanged, since the application of the remapping functions also maintains the color ratios.

They use the HSV color space [21] and modify the V channel (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.

2.2.2 Quality Metrics

Following the initial exposure fusion, the blending weights are computed from quality metrics. However, only the well-exposedness measure is kept from the Mertens et al. method, which uses two more measures based on contrast and saturation. Indeed, the saturation cannot be used, since the generated images have only one luminance channel. As for the contrast metric, it is replaced by a direct deduction of the contrast from the remapping functions.

Well-exposedness. The *well-exposedness* measures how close each pixel's value is to the median value 0.5 using a Gauss curve

$$w_{e,k}(\mathbf{x}) = \exp \frac{-\left(\widehat{u}_k(\mathbf{x}) - 0.5\right)^2}{2\sigma^2},\tag{8}$$

where $\sigma = 0.2$.

Contrast. While exposure fusion's *contrast metric* uses the absolute value of a discrete Laplacian filter (the sum of differences over the four nearest neighbors) applied to the grayscale version of the image, in the simulated exposure fusion [8] the contrast metric is directly deduced from the remapping functions, by using their derivative. This is also more efficient. The metric is computed using

$$w_{c,k}(\mathbf{x}) = \begin{cases} \alpha^{|k|/N_{\max}}(g' \circ f^*)(u_k(\mathbf{x}), k), & \text{if } k < 0\\ \alpha^{k/N_{\max}}(g' \circ f)(u_k(\mathbf{x}), k), & \text{otherwise,} \end{cases}$$
(9)

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

$$g'(t,k) = \begin{cases} 1, & \text{if } |t - \rho(k)| \le \frac{\beta}{2} \\ \frac{\lambda^2}{(|t - \rho(k)| - b)^2}, & \text{otherwise.} \end{cases}$$
(10)

An example of the remapping functions with their derivatives is presented in Figure 3.



Figure 3: Remapping functions and their derivatives for $\alpha = 4$, $\beta = 0.6$. The restrained dynamic range location is adapted to the different generated images, so as to preserve the most pertinent information. The number of under- and over-exposed images itself is automatically computed from the histogram.

2.2.3 Weight Maps

Like EF, SEF computes the pixel-wise weight maps by multiplying the quality metrics, then normalizing them. The power functions added in EF to allow the user to give different degrees of importance to the different measures are removed in SEF however, due to their little influence on the final result.

The final normalized weight maps are then obtained as

$$\widehat{w}_{k}(\mathbf{x}) = \frac{w_{e,k}(\mathbf{x}) \cdot w_{c,k}(\mathbf{x})}{\sum_{k} w_{e,k}(\mathbf{x}) \cdot w_{c,k}(\mathbf{x})}.$$
(11)

2.2.4 Multiscale Fusion

The fusion of the generated sequence is then carried out using the Ogden et al. multiscale fusion [16]. 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 shall denote by $L\{u\}$ the Laplacian pyramid of the input image u, by $G\{w\}_k$ the Gaussian pyramid of the weights w_k , and by l the scale. The blending operation is then

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

$$(12)$$

We refer to [7] for a precise description, with pseudo-codes, of the multiscale blending procedure.

Low-frequency halo and deeper pyramids The restrained dynamic range strategy also permits to use deeper pyramids for the fusion. Indeed, the Mertens et al. method does not use the deepest possible pyramid. The reason is that a deeper pyramid increases the out-of-range artifact. However, by stopping the blending slightly earlier than when the residual's size is just one pixel-wide, some low-frequency halos can become visible. In addition to proposing a solution to the out-of-range artifact, the authors of [8] suggest to use a deeper pyramid to remove the second known artifact of EF, its low-frequency halo [15]. The low-frequency halo does not affect the simulated exposure fusion [8], so we do not pursue the discussion on this artifact further.

2.2.5 Robust Normalization

After the multiscale fusion, the color is re-introduced. This operation boils down to replacing the v channel of the input image expressed in the HSV color space by the enhanced luminance ℓ' , and to convert back to the RGB color space.

This operation performed, the fused image's intensities need to be rescaled to the unit range. Depending on β and the content of the image, this most generally amounts to stretch the intensity range (as opposed to EF, which typically compresses this range). This final step over-stretches the colors to [0, 1] by allowing 1% of clipping on both sides of the histogram. Instead of treating the channels separately [12], they are handled together to preserve colors. This operation is called Robust Normalization. We refer to [7, 12] for an implementation.

Algorithm 2: Simulated Exposure Fusion (SEF)								
	input : image u and parameters α , β							
	output : enhanced image v							
1	$\ell \leftarrow v$ channel of input in HSV color space							
2	$M, N, N^* \leftarrow \text{SEQUENCELENGTH}(\alpha, \beta, \text{median}\{u\})$	// Algorithm 1						
3	for $k \in \{-N^*, \dots, N\}$ do	<pre>// Generate a sequence</pre>						
4	if $k < 0$ then							
5	$ \left \begin{array}{c} \widehat{\ell}_k \leftarrow (g \circ f^*)(\ell, k) \end{array} \right $	// under-exposed image, Equation $\left(1 ight)$						
6	$ w_{c,k} \leftarrow \alpha^{ k /N_{\max}}(g' \circ f^*)(\widehat{\ell}_k, k) $	// contrast weights, Equation $\left(9 ight)$						
7	else							
8	$\widehat{\ell}_k \leftarrow (g \circ f)(\ell, k)$	// over-exposed image, Equation $\left(1 ight)$						
9	$ w_{c,k} \leftarrow \alpha^{k/N_{\max}}(g' \circ f)(\widehat{\ell}_k, k) $	// contrast weights, Equation $\left(9 ight)$						
10	$w_{e,k} \leftarrow \exp\left(-(\widehat{\ell}_k - 0.5)^2 / (2\sigma^2)\right) $ //	well-exposedness weights, Equation (8)						
11	$\widehat{w}_{k} \leftarrow \left(w_{e,k} \cdot w_{c,k}\right) / \left(\sum_{k} w_{e,k} \cdot w_{c,k}\right)$	// final, normalized weights, Eq. (11)						
12	$\ell' \leftarrow \text{MultiscaleFusion}(\widehat{\ell}, \widehat{w})$	// Equation (12)						
13 $v_{\text{RGB}} \leftarrow \text{convert to RGB}$ the input in HSV color space where V is replaced by ℓ'								
14 $v_{\text{rgb}} \leftarrow \text{RobustNormalization}(v_{\text{rgb}})$								

3 Results and Comparisons

Image contrast enhancement techniques belong to a wider class of methods, the tone-mapping operators. 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. Yet the very same methods adapt to LDR images and improve their brightness and contrast. Among them, *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 tone mapping methods introduced for LDR images we can mention Automatic Color Enhancement (ACE) [4, 5], Local Color Correction (LCC) [20, 6, 13], the Fast Local Laplacian Filter (LLF) [17, 1]. Contrast enhancement techniques take inspiration in Land's Retinex method [10] which modeled the human color perception system. Several different implementations of the Retinex method exist. The most popular is arguably Multi-Scale Retinex (MSR) [18, 9].

In this section, we compare the results of SEF [8] with other Retinex and histogram-based methods: ACE, MSR, Histogram Equalization (HE), Contrast-Limited Adaptive Histogram Equalization (CLAHE) [19], LCC [20, 6], LLF [17, 1] and two other single-image contrast-enhancement methods derived from exposure fusion [14] like SEF: the Detail-Enhanced Exposure Fusion (DEEF) of Li et al. [11], and the Backlit Image Enhancement algorithm (BIE) of Buades et al. [2]. In the DEEF method, images are simulated and fused with a modified version of [14], the exposure fusion of Mertens et al. The authors aim at enhancing backlit images by fusing three images with increased contrast. The first simulated image is obtained from the input using a tone-curve increasing the contrast in the dark parts, and leaving the bright regions untouched. The two others are obtained by applying factors superior to one to this first simulated image. Likewise, the BIE method is intended for the enhancement of backlit images. Ten images are created from the initial one using global mappings. Five images are obtained using gamma functions and five using logarithmic functions. The gamma functions use exponents superiors to one so the contrast in the bright areas is enhanced. Then a multiscale fusion method is used to fuse these eleven images (including the input one). The authors use Mertens' method, with two modifications: (1) the fusion is applied per channel, thus the saturation metric is discarded; (2) the contrast metric is replaced by a local variance. After the fusion, a sharpening method is used for local contrast enhancement. These operations are applied separately on the channels, which are then recombined. The final operation is a chrominance correction which restores the original R/G/B ratios. These two methods do not use restrained ranges like SEF. Results are shown in Figure 4.

The results of SEF are in general clearer than the other methods. In the "Columns" image for example, the SEF and BIE results are the only ones that really enhance the bushes and the house wall between the columns. Furthermore, the SEF method does not create halos, contrarily to BIE and CLAHE (and also, to a lesser extent MSR and LLF). The halos in BIE [2] probably originate from the contrast enhancement step. The halos of this method are particularly visible in the Lighthouse and Airplane images. The local contrast however is not better in BIE than in SEF: compare for example the trees on the right of the Girl image. SEF does not lose contrast anywhere in the image, contrarily to HE, ACE, DEEF and MSR. Unlike CLAHE, SEF preserves the global contrast: for example, the sky of the "Columns" image is not darkened. Also, SEF does not over-enhance some regions that do not need it, contrarily to LLF in the "Es Canonge" image for example. LLF enhances the already correctly exposed parts of the image, resulting in too much local contrast on the sky, or in the hair in the "Girl" image. DEEF does not enhance bright regions. Besides, it produces images with values largely outside [0, 1] (the out-of-range artifact), thus needs compression. This means that bright regions systematically lose contrast. This effect is visible in the sky of the "Columns" and "Es Canonge" images. Moreover, the dark parts in DEEF are not as well enhanced as in SEF. LCC gives grayish images with low contrast.

In short, SEF gives as good results as state-of-the-art Retinex algorithms such as MSR, ACE, and CLAHE, but better preserves and enhances information everywhere without creating halos. Furthermore, the results look more natural. The DEEF method suffers from the out-of-range artifact, which produces images with very poor local contrast. Besides, Li et al. 's [11] DEEF is more complex and costly since it requires edge-aware smoothing in addition to a Laplacian pyramid blending. The BIE method suffers from halos, and, like DEEF, is more costly than SEF, since eleven images are systematically created and fused (whereas SEF generally uses around five), and supplementary steps are added for contrast enhancement and chrominance correction. Concerning edge-aware local contrast enhancement, the state-of-the-art LLF is able to correctly enhance contrast everywhere without creating artifacts (slight halos are nevertheless sometimes visible). Compared to LLF, SEF



Figure 4: Comparison of SEF with other enhancement methods on five 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 $N^* = 0$. For LLF, we used the default parameters $\sigma = 0.1$, N = 10, which means that 10 images were fused. For LCC, ACE and MSR we used the default parameters proposed in [6, 5, 18]. For BIE and DEEF we used the authors' implementation default parameters. For HE and CLAHE we used Matlab implementation's default parameters. (The reader is invited to zoom in the pdf for a more accurate comparison).



Figure 5: Comparison of the fused results for several values of the parameters $\alpha \in [2, 16]$ and $\beta \in [0.2, 0.8]$, which are the typical ranges. The values $\alpha = 8$ and $\beta = 0.5$ are recommended by the authors in [8]. The fused images are good and artifact-free for a wide range of values. Nonetheless, the combination of small α and small β seems not desirable, as the contrast of the dark areas is not sufficiently enhanced. Under each image: the number of under- and over-exposed images, respectively, and the factor applied in the final robust normalization step.

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. Indeed, for the results displayed in Figure 4, SEF only required the fusion of three Laplacian pyramids, while LLF used ten.

3.1 Effect of the Parameters

In Figure 5, we applied SEF to one image but with different values of the two parameters α and β . The original image is displayed in Figure 4, as well as its enhanced version with the values recommended in [8]: $\alpha = 8$ and $\beta = 0.5$. In this experiment, we explore SEF's result in the range of realistic values for its two parameters, together with their interactions. The acceptable ranges seem to be [2, 16] for the contrast enhancement parameter α , and [0.2, 0.8] for the restrained dynamic of the generated images, β .

We observe that the enhanced images are good in a wide range of values. In fact, the results are artifact-free whatever the parameters' values. However, some results are better than others. In particular, the column $\alpha = 2$ has the poorest results, because the enhancement is insufficient.

Although they remain acceptable for large β , because the produced image is close to the input, the enhanced image with both small α and β is not visually pleasing, mostly due to the lack of local contrast in the originally dark parts.

Under the enhanced images are also displayed the number of under- and over-exposed simulated images N^* and N, respectively, together with the final factor applied to the intensities at the robust normalization step. Observe that the combination of large α and β is not recommendable either, since it leads to factors inferior to 1, meaning that some contrast is lost.

To sum up, the best results are obtained for intermediate values of β , i.e. between 0.4 and 0.6. Furthermore, we notice a dependency between α and β : small contrast enhancement values α need larger restrained ranges β , and, conversely, large contrast enhancement factors need smaller restrained ranges. That being said, the SEF method is capable of extreme contrast enhancement without artifact, e.g. with $(\alpha, \beta) = (16, 0.4)$ and (16, 0.2).

4 Conclusion

We have described and analyzed Simulated Exposure Fusion, introduced in [8] to locally enhance the contrast in an image. This method makes the most of exposure fusion [14, 15], a well known and widely used fusion method since its publication in 2007. The two main contributions of SEF are: a generation model capable of simulating a sequence from a single input; a correction of EF that resolves its out-of-range artifact. SEF is capable of extreme enhancement while maintaining a natural rendition. Furthermore, it automatically adapts to the global illumination of the input, hence targeting the enhancement efforts only on areas that require it. No artifacts were observed in the results. Following our analysis of the parameters' impact, we recommend a range compression parameter β in [0.4, 0.6] and to avoid the use of small enhancement factors α when β is small.

Image Credits



Acknowledgment

We thank Jean-Michel Morel for fruitful comments and discussions. Work partly financed by Office of Naval research grant N00014-17-1-2552, DGA Astrid project "filmer la Terre" n°ANR-17-ASTR-0013-01, and ANRT CIFRE Ph.D. scholarship n°2014/1323 of the French Ministry for Higher Studies, Research and Innovation.

References

[1] M. AUBRY, S. PARIS, S. W. HASINOFF, J. KAUTZ, AND F. DURAND, *Fast and robust pyramid-based image processing*, research report, Computer Science and Artificial Intelligence Laboratory, November 2011.

- [2] A. BUADES, J.-L. LISANI, A. B. PETRO, AND C. SBERT, Backlit images enhancement using global tone mappings and image fusion, IET Image Processing, (2019). https://doi.org/10. 1049/iet-ipr.2019.0814.
- [3] P. BURT AND E. ADELSON, The Laplacian pyramid as a compact image code, IEEE Transactions on Communications, 31 (1983), pp. 532-540. https://doi.org/10.1109/TCOM.1983. 1095851.
- [4] C. GATTA, A. RIZZI, AND D. MARINI, ACE: An automatic color equalization algorithm, in Conference on Colour in Graphics, Imaging, and Vision, Society for Imaging Science and Technology, 2002, pp. 316–320.
- [5] P. GETREUER, Automatic Color Enhancement (ACE) and its Fast Implementation, Image Processing On Line, 2 (2012), pp. 266-277. https://doi.org/10.5201/ipol.2012.g-ace.
- [6] J.G. GOMILA SALAS AND J.L. LISANI, Local Color Correction, Image Processing On Line, 1 (2011), pp. 260–280. https://doi.org/10.5201/ipol.2011.gl_lcc.
- [7] C. HESSEL, An Implementation of the Exposure Fusion Algorithm, Image Processing On Line, 8 (2018), pp. 369–387. https://doi.org/10.5201/ipol.2018.230.
- [8] C. HESSEL AND J.-M. MOREL, An extended exposure fusion and its application to single image contrast enhancement, in IEEE Winter Conference on Applications of Computer Vision (WACV), IEEE, 2020.
- [9] 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 (1997), pp. 965–976. https://doi.org/10.1109/83.597272.
- [10] E. H. LAND, The retinex theory of color vision, Scientific American, (1977), pp. 108–129.
- [11] Z. LI, Z. WEI, C. WEN, AND J. ZHENG, Detail-enhanced multi-scale exposure fusion, IEEE Transactions on Image processing, 26 (2017), pp. 1243–1252. https://doi.org/10.1109/TIP. 2017.2651366.
- [12] N. LIMARE, J.L. LISANI, J-M. MOREL, A.B. PETRO, AND C. SBERT, Simplest Color Balance, Image Processing On Line, 1 (2011), pp. 297–315. https://doi.org/10.5201/ipol. 2011.llmps-scb.
- [13] J. L. LISANI, Adaptive local image enhancement based on logarithmic mappings, in IEEE International Conference on Image Processing (ICIP), IEEE, 2018. https://doi.org/10.1109/ ICIP.2018.8451655.
- [14] T. MERTENS, J. KAUTZ, AND F. VAN REETH, Exposure fusion, in Pacific Conference on Computer Graphics and Applications, 2007. https://doi.org/10.1109/PG.2007.17.
- [15] —, Exposure fusion: A simple and practical alternative to high dynamic range photography, Computer Graphics Forum, 28 (2009), pp. 161–171. https://doi.org/10.1111/j.1467-8659. 2008.01171.x.
- [16] J. M. OGDEN, E. H. ADELSON, J. R. BERGEN, AND P. J. BURT, Pyramid-based computer graphics, RCA Engineer, 30 (1985), pp. 4–15.

- [17] S. PARIS, S. W. HASINOFF, AND J. KAUTZ, Local Laplacian filters: edge-aware image processing with a laplacian pyramid, ACM Transactions on Graphics, 30 (2011), p. 68.
- [18] A.B. PETRO, C. SBERT, AND J-M. MOREL, Multiscale Retinex, Image Processing On Line, (2014), pp. 71–88. https://doi.org/10.5201/ipol.2014.107.
- [19] S. M. PIZER, E. P. AMBURN, J. D. AUSTIN, R. CROMARTIE, T. GREER A. GESELOWITZ, B. TER HAAR ROMENY, J. B. ZIMMERMAN, AND K. ZUIDERVELD, *Adaptive histogram equalization and its variations*, Computer Vision, Graphics, and Image Processing, 39 (1987), pp. 355– 368. https://doi.org/10.1016/S0734-189X(87)80186-X.
- [20] A. RIZZI, C. GATTA, AND D. MARINI, A new algorithm for unsupervised global and local color correction, Pattern Recognition Letters, 24 (2003), pp. 1663–1677. https://doi.org/10.1016/ S0167-8655(02)00323-9.
- [21] A. R. SMITH, Color gamut transform pairs, ACM SIGGRAPH Computer Graphics, 12 (1978), pp. 12–19.