A Retinex Inspired Bilateral Filter for Enhancing Images under Difficult Light Conditions

Michela Lecca[a]
Fondazione Bruno Kessler, via Sommarive 18, Trento, Italy

Keywords: Retinex, Bilateral Filter, Spatial Color Algorithm, Image Enhancement.

Abstract: This paper presents SuPeR-B, a novel, Retinex inspired color spatial algorithm to enhance images acquired under difficult light conditions, such as pictures containing dark and bright regions caused by backlight and/or local, not diffused spotlight. SuPeR-B takes as input a color image and improves its readability by processing its color channels independently in accordance with some principles of the Retinex theory. Precisely, SuPeR-B re-works the channel intensity of each pixels accounting for differences computed both in the spatial and intensity domains. In this way, SuPeR-B acts as a bilater filter. The experiments, carried out on a real-world dataset, shows that SuPeR-B ensures good enhancement results, also in comparison with other state-of-the-art algorithms: SuPeR-B improves the overall content of the image, making the dark regions brighter and more contrasted, while lowering possible chromatic dominants of the light.

1 INTRODUCTION

An image enhancer is an algorithm that improves the visibility of the content and of the details of an input image. Such an algorithm is particularly needed each time an image has been acquired under difficult light conditions, like low-light, back-light and/or multiple light sources, which cause noise, color distortions and strong shadows and make hard to understand what the image depicts. Improving the image quality is essential not only for human observers but also for all the machine vision algorithms requiring an accurate detail visibility, such as key-point detectors for robust and illumination invariant image/object recognition (Lecca et al., 2019). Many methods for enhancing pictures captured under bad illuminations have been proposed in the literature. Some examples are global and local image methods exploiting statistical analysis, e.g. (Zuiderveld, 1994), (Gianini et al., 2014), Retinex inspired algorithms and variants, e.g. (Morel et al., 2010), (Banić and Lončarić, 2013), (Lecca et al., 2018), (Lecca and Messelodi, 2019), (Banić and Lončarić, 2015), illuminance/reflectance decomposition approaches, e.g. (Guo et al., 2017), (Wang et al., 2013), (Fu et al., 2016), machine learning techniques, e.g. (Jiang et al., 2019), (Wei et al., 2018), (Lv et al., 2018), (Li et al., 2020).

Despite the huge efforts made till now, image enhancement is still an open issue. In fact, the most enhancers rely on specific hypotheses about illumination, reflectance and spectral properties of the acquisition device. These assumptions limit the enhancer applicability to specific contexts. For instance, many algorithms assume that the illumination varies slightly across the image, but this prevents the enhancement of images with abrupt changes of light, as for instance strong shadows. In this respect, the enhancement of images with very dark and very bright regions has been poorly investigated. These images are usually generated by capturing scenes with backlight or local, brilliant but not diffused spotlights (see fig. 1).

This work contributes to the state of the art on the enhancement of such images by presenting SuPeR-B, a novel spatial color algorithm obtained as a variant of the Retinex inspired image enhancer SuPeR (Lecca and Messelodi, 2019). Both SuPeR and SuPeR-B take as input a color image, process their channels separately, and partition each channel by regular, not overlapping tiles, which are treated as *su*per-*pi*x*e*ls and processed according to some principles of the *R*etinex theory. The acronym SuPeR comes from Italic letters marked above. SuPeR rescales the intensity $I(x)$ of any pixel x by a value, which is computed as the average of the tile maximum intensities greater than $I(x)$, where each intensity is weighted by a function inversely proportional to the distance

[a] https://orcid.org/0000-0001-7961-0212

Lecca, M.
A Retinex Inspired Bilateral Filter for Enhancing Images under Difficult Light Conditions.
DOI: 10.5220/0010235900760086
In *Proceedings of the 16th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP 2021) - Volume 4: VISAPP*, pages 76-86
ISBN: 978-989-758-488-6

Figure 1: Examples of images captured with backlight and with local, not diffuse spotlight and their versions enhanced by SuPeR-B (α = 0, a, b = -1). In the backlight image (top, left), the subject is displayed against a brilliant sky and appears dark. In the spotlight image (top, right), the monitor of the notebook positioned in a dark room acts as a local, not diffused spotlight, that cannot illuminate sufficiently the near-by regions. In both the cases, the image content without enhancement is unreadable.

of the tile barycenter from x. The resulting image is an enhanced version of the input one: it is brighter and more contrasted, has a lower color distribution entropy, while shadows and color casts possibly due to the light are smoothed or even removed. The experiments show that SuPeR provides an accurate image enhancement of the most of real-world, low-light images with very fast computational times, but its performance decreases in case of backlight and spotlight images. SuPeR-B (where 'B' stands for 'back-light' but also indicates that this is a second, improved version of SuPeR) overcomes this problem by implementing a novel function weighting the tile intensities involved in the enhancement of $I(x)$. This function, modeled by a Coon patch, accounts not only for the distance between the tile barycenters and x but also for the amount of the difference between the tile maximum intensity and $I(x)$. In this way, SuPeR-B acts like a bilateral filter which processes pixelwise the image based both on intensity and spatial features, brightening dark regions within edge preservation. The experiments carried out on a dataset of real-world back-light and spotlight images show that SuPeR-B outperforms SuPeR as well as other algorithms at the state of the art, in particular the multi-scale Retinex algorithm (Petro et al., 2014), the channel-division approach (Ramirez Rivera et al., 2012), the image fusion-based enhancer for single backlit images (Wang et al., 2016) and the learning based approach for image restoration (Li and Wu, 2018). In this framework, the performance of SuPeR-B have been evaluated by comparing the brightness, the contrast and the color distribution entropy of the test images before and after the enhancement. The results show that SuPeR-B increases the brightness and contrast, lowers the color distribution entropy and preserves the important edges, meaning that the enhanced image has more visible details and a more readable content than the input one.

2 RELATED WORK

The enhancement of images with dark and bright regions due to the light has been scarcely addressed in the literature. From the hardware point of view, the high dynamic range (HDR) devices attempt to deal with the root of the problem, by capturing multi-exposure images of the same scene and merging them via tone mapping functions in order to generate a high quality pictures. Nevertheless, the HDR tone mapping functions often introduce in the final image unpleasant distortions and artifacts, and thus they themselves are subject of current research. In addition, HDR imaging cannot solve the problem of enhancing an existing image where the light conditions hamper the visibility of the content of some regions.

The image enhancers working globally, such as histogram equalization, generally perform scarcely because they do not account for the variability of the light across the scene. Better performance is reached by spatial adaptive enhancers, such as many Retinex inspired algorithms. These latter basically process the image channels separately and re-work the channel intensity of each pixel x based on the intensity (and sometimes also on other visual features like gradient) of pixels sampled from a neighborhood of x. The output image is brighter and more contrasted than the input one, possible shadows and color casts due to the illumination are smoothed or even removed, the dynamic range of the image is stretched. In other Retinex inspired algorithms, the local spatial information is modeled through a function of the distance of x from the pixels selected as relevant for the enhancement, e.g. (Jobson et al., 1997), (M. Lecca, A. Rizzi, and R.P. Serapioni, 2017), (Lecca, 2018), (Lecca and Messelodi, 2019). For all the Retinex inspired methods, the definition of the *locality*, i.e. of the neighborhood of x to be processed or of the distance function, is a key point. For instance, when the all the pixels around x are considered, these algorithms tends to behave like the Max-RGB algorithm. This latter rescales the intensity of each pixel by the maximum intensity over the image, but this leads to unsatisfactory results on most real world and noisy images. When the neighborhood of x is very small, the final image is close to an edge map, where chromatic information and some details are lost. A priori knowledge about image content and/or light sources may greatly help to choose the neighborhood or distance function maximizing the algorithm performance, but such a knowledge is often unavailable.

Retinex algorithms at multiple resolution allow to simulate the effects of HDR imaging (Jobson et al., 1997), (Petro et al., 2014): the input image is processed sequentially by a Retinex algorithm with increasing neighborhood (or support for the distance function) and the obtained images are averaged together. This approach avoids the user to select a specific extent of the pixel neighborhood while stretches the dynamic range of the image, brightening the dark regions. Nevertheless, the prize to pay is a longer computational time, a less accurate removal of the light effects and sometimes the generation of artifacts.

Other algorithms, e.g. (Tsai and Yeh, 2010), (Li and Wu, 2018), (Ramirez Rivera et al., 2012), specifically address the problem of enhancing backlight/spotlight images. To this purpose, they segment the input image in backlight/frontlight regions and rework them with different enhancing functions. In particular, the work in (Tsai and Yeh, 2010) segments the image by thresholding, then linearly stretches and shifts separately the dark and bright regions in order to brighten the first one and darken the second one. A variant of this computational scheme is proposed by (Li and Wu, 2018) (hereafter denoted as Backlit), that replaces the thresholding based segmentation with a region growing approach and introduces guided filters for improving the results. Specifically, the image is partitioned in backlight and frontlight regions through support vector machines and conditional random fields, and each segmented region is enhanced by a tone mapping function which maximizes the Weber contrast while minimizes tone distortions. Region borders are processed by a linear combination of the tone mapping functions estimated before to avoid artifacts. For all these approaches, the use of different functions for enhancing the dark and bright regions enables good results, which however strongly depend on the segmentation accuracy. The work in (Ramirez Rivera et al., 2012), here called Channel division, enhances edges and flat regions with different approaches considering image texture. The enhancement results obtained on edges and flat regions are blended together to highlight details while maintaining the smoothness of flat regions. This method provides good results for a wide range of images, but it poorly works in extreme conditions, as for instance on pictures with near-black portions, where the contrast computation is affected by noise. The work in (Wang et al., 2016), here shortly named Fusion, proposes a fusion method, that processes the input image I in the HVS space and computes from I an image I_1 with over-enhanced dark regions, an image I_2 where the dynamic range of bright regions has been compressed, and image I_3 with enhanced contrast. Each image I_i is smoothed by a Laplacian operator to remove undesired halos, then it is pixel-wise multiplied to a weight that controls its overall exposure. The resulting images are averaged and the result is the enhanced version of I.

The difficulty in defining the spatial locality of the image processing, the generation of artifacts along the edges, the low quality of the image signal in dark regions, the lack of a priori knowledge about the image content are issues that make hard and challenging the enhancement of a bad illuminated image and this justifies the research of new methods.

3 THE PROPOSED APPROACH

This Section presents SuPeR (3.1) and its variant SuPeR-B (Subsection 3.2).

3.1 SuPeR

Among the many Retinex implementations and in particular among the algorithms of the Milano Retinex family (Rizzi and Bonanomi, 2017), the Retinex inspired spatial color algorithm SuPeR is of interest not only for its enhancement performance, but also for its fast computational times, requiring less than 0.2 seconds to process an image with size 640×480 with 100 tiles on a standard PC with CPU Intel(R) Xeon(R) CPU E3-1245 v6 at 3.70 GHz (see (Lecca et al., 2019)).

SuPeR takes as input a color image $\overline{I}$ and an integer parameter $n > 0$. It enhances $\overline{I}$ by implementing some principles which are at the basis of the Retinex theory, i.e. 1) independent processing of the three image channels; 2) channel processing based on local spatial and intensity information; 3) image enhancement with smoothing/removal of light effects.

Let us introduce some notation. Let I be a channel of $\overline{I}$ and let S be the domain of $\overline{I}$, i.e. the set of the spatial coordinates of the pixels composing $\overline{I}$ (and thus I). Let D denote the length of the diagonal of S. Here, the channel I is regarded as a function from S to the intensity range (0, 1], where zero has been excluded to enable the computation of intensity ratios. Let L indicate the enhanced version of I.
SuPeR enhances $\overline{I}$ by implementing sequentially a global and a pixel-wise image processing.

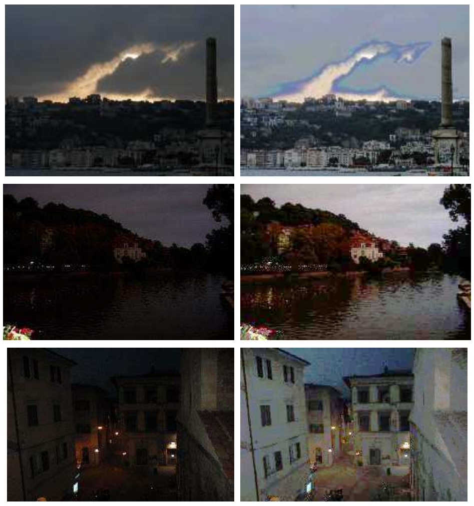

Figure 2: Examples of images (on left) enhanced by SuPeR (on right) with n = 25.

In the global processing, SuPeR partitions S in n rectangular, not overlapping tiles $T_1, \dots, T_n$ by a regular grid superimposed on S. According to the principle 1 reported above, SuPeR processes each channel independently. Specifically, given the channel I, SuPeR computes the set $B = \{(b_i, m_i) : i = 1, \dots, n\}$, where $b_i \in S$ is the barycenter of T_i and m_i is the maximum channel intensity over T_i.
In the pixel-wise processing, SuPeR maps the intensity $I(x)$ of any $x \in S$ on a new value $L(x)$ defined as follows:

$$L(x) = \begin{cases} \frac{\sum_{(b,m)\in\mathcal{B}_x}(1-d(x,b))\frac{I(x)}{m}}{\sum_{(b,m)\in\mathcal{B}_x}(1-d(x,b))} & \text{if } \mathcal{B}_x \neq \emptyset \\ 1 & \text{otherwise} \end{cases} \quad (1)$$

where $\mathcal{B}_x = \{(b,m) \in B : I(x) < m\}$ and d is the Euclidean distance between x and b, normalized to range between 0 and 1:

$$d(x,b) = \frac{\| x - b \|^2}{D^2}. \quad (2)$$

According to Equation (1), given the point x, SuPeR selects from B the subset $\mathcal{B}_x$ composed by the pairs (b,m) corresponding to tiles whose maximum intensity exceeds $I(x)$. If $\mathcal{B}_x$ is empty, then the intensity $I(x)$ is mapped onto one, otherwise it is rescaled by the values ms and each ratio $I(x)/m$ is averaged with a weight d depending on the distance of b from x. This weight models the spatial locality of the algorithm and decreases by increasing the distance of b from x. This behaviour is in line with the Retinex principle stating that the intensities close to x influence the perception of $I(x)$ more than those located further (principle 2).

The final enhanced image is obtained by packing the L's into an RGB image, which is robust to changes of light (principle 3). In fact, the division of $I(x)$ by other intensity values enables smoothing or even discounting light effects, such as shadows or color dominants of the illumination. This is in line with the von Kries model (Finlayson et al., 1994), (Lecca, 2014), stating that the change of any pixel RGB triplet caused by a light variation can be approximated by a linear diagonal transform of the RGB triplet. In the original paper of SuPeR, the authors suggest to smooth the input image by a median filter so that to make the method robust to salt and pepper noise.

The experiments on SuPeR, conducted on real-world images, showed good performance, both in terms of enhancement and of computational time. The images were improved by increasing their brightness and their contrasts and by flattening their color distribution. Some examples are shown in Fig. 2. Nevertheless, in case of images with extreme conditions, such as backlight and spotlight, the performance of SuPeR decreases, as illustrated in Fig. 3 and discussed in the next Subsection.

3.2 SuPeR-B

Figure 3 shows an example of image with extreme backlight, depicting a giraffe against a brilliant, reddish sky. The enhancement by SuPeR smooths the effects of the light changing the colors of the sky, but the giraffe appears still dark and its details are undetectable. Regarding the results of the other algorithms specifically designed to cope with backlight, the algorithm Channel Division (Ramirez Rivera et al., 2012) performs poorly, while the others (i.e. the multi-scale Retinex MSR described in (Petro et al., 2014), Fusion (Wang et al., 2016) and Backlit (Li and Wu, 2018)) work a little better, but the visibility of the giraffe remains low.

The bad result by SuPeR is due to the fact that the intensity values of the pixels inside the dark regions are divided by much greater intensity values sampled from the sky. Although penalized by the distance function, these sky intensities contribute heavily to the image enhancement and make the values of L decreasing on the giraffe region. As already mentioned in Section 3, multi-resolution Retinex inspired algorithms, as MSR, tend to generate artifacts and they do not ensure a good removal/smoothing of the light effects. This discourages the application of SuPeR at multiple scales. Thresholding the distance function in Equation (1) to exclude the high sky intensities from the enhancement of the dark region may appear a possible way to increase the giraffe intensities, but it may cause the loss of important edges, as those on the borders. A novel solution is proposed by SuPeR-B.

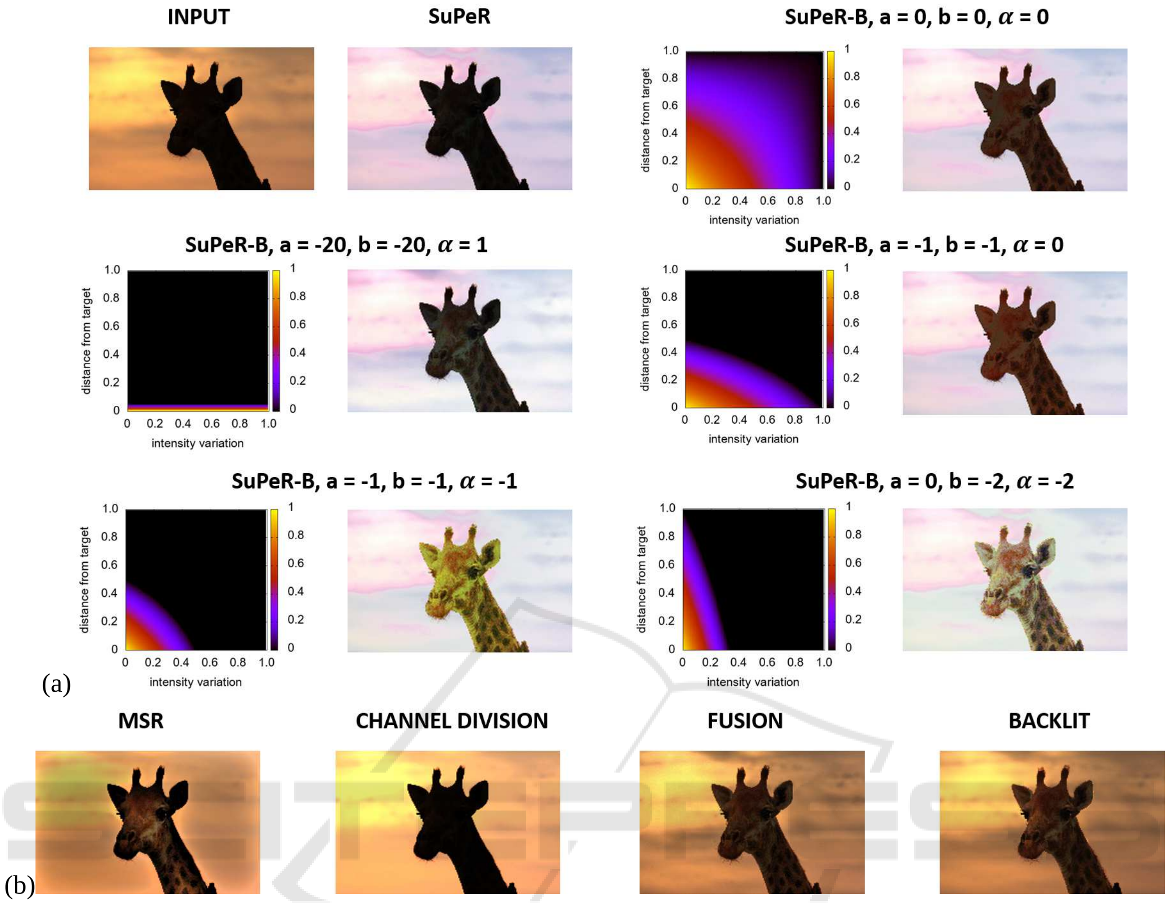

Figure 3: (a) Image enhancement of a backlight image by SuPeR and by SuPeR-B for different values of the parameters a, b and α. For SuPeR-B the weighting function is also shown. (b) Image enhancement of the input image in (a) by the algorithms compared with SuPeR-B.

SuPeR-B inherits from SuPeR the general workflow, i.e. it processes the image channels separately by the two global and local routines of SuPeR, but introduces in Equation (1) a novel function f that weights the contribution of the intensities in $\mathcal{B}_x$ not only upon their spatial distance from x but also upon their difference from $I(x)$. Specifically, SuPeR-B replaces the equation (1) with the following one:

$$L_B(x) = \begin{cases} \frac{\sum_{(b,m)\in\mathcal{B}_x} f(\delta I(x,b), d(x,b)) \frac{I(x)}{m}}{\sum_{(b,m)\in\mathcal{B}_x} f(\delta I(x,b), d(x,b))} & \text{if } \mathcal{B}_x \neq \emptyset \\ 1 & \text{otherwise} \end{cases} \quad (3)$$

where $\delta I(x,b) = I(b) - I(x)$ and f weights the intensities close to x and to $I(x)$ more than the other values. While on bright regions the values of f is almost irrelevant, on dark regions they are set to increase the brightness and the detail visibility.

There exist many expressions for f and different expressions produce different enhancement levels. In the current implementation of SuPeR-B, f is modeled by a Coon patch, i.e. a compact surface in 3D space whose borders are described by paths in 2D space intersecting two by two at the four patch corners (see Figure 4 for an example). Precisely, let $c_0, d_0, c_1, d_1 : [0,1] \rightarrow \mathbf{R}$ be four continuous paths with $c_0(0) = d_0(0)$, $c_0(1) = d_1(0)$, $c_1(0) = d_0(1)$ and $c_1(1) = d_1(1)$. The equation of the Coon patch bounded by these paths is:

$$C(s,t) = S(s,t) + T(s,t) - U(s,t) \quad (4)$$

where

$$\begin{aligned} S(s,t) &= (1-t)c_0(s) + tc_1(s) \\ T(s,t) &= (1-s)d_0(t) + sd_1(t) \\ U(s,t) &= c_0(0)(1-s)(1-t) + c_0(1)s(1-t) + \\ &\quad c_1(0)(1-s)t + c_1(1)st. \end{aligned}$$

Pictorially, we can imagine that the endpoints of the path c_0, which belongs to the xz plane, *move* respectively along the paths d_0 and d_1 modifying the shape of c_0 accordingly until c_0 *rests* (and coincides) with c_1. The same representation holds for d_0, whose endpoints are initially over $c_0(0)$ and $c_1(0)$ and go respectively to $d_1(0)$ and $d_1(1)$ until d_0 reaches (and coincide with) d_1. Moving the corners and/or changing the path equations allows to model a lot of surfaces.

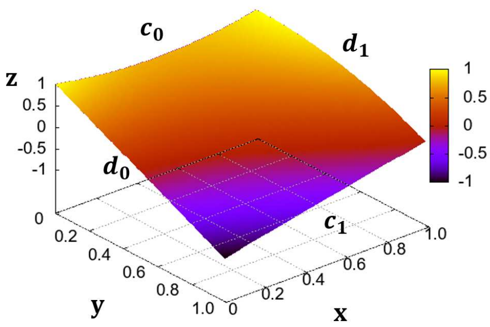

Figure 4: Example of Coon surface in 3D space.

In the current version of SuPeR-B, the paths c_0, c_1, d_0 and d_1 are lines defined as follows:

$$\begin{aligned} c_0(s) &= s(\alpha-1)+1 \\ c_1(s) &= s(b-a)+a \\ d_0(t) &= t(a-1)+1 \\ d_1(t) &= t(b-\alpha)+\alpha \end{aligned}$$

where α, b, a are real-world user parameters and $s,t \in [0,1]$. In order to ensure positivity of the weights, the function f is defined as

$$f(s,t) = \max(C(s,t),0), \tag{5}$$

where the parameters s and t represent respectively the variation of intensity δI and the value of d between two image pixels.

The values of α, a, b must be chosen so that (i) to satisfy the Retinex principle stating that the contribution of the sampled pixels to the enhancement decreases with their distance from x, and (ii) to improve the visibility of the image content in the dark regions by weighting more the items of $\mathcal{B}_x$ closest to $I(x)$ in the intensity space. As a general condition, for α, a, b such that

$$\alpha \leq 1, a < 1, b \leq \min\{a, \alpha\} \tag{6}$$

the requirements (i) and (ii) are satisfied. In fact, when $a \geq 1$, the intensities spatially located far from x become more relevant than those close to x in the computation of $L(x)$ and this violates the requirement (i). On the contrary, values of $a < 1$ enables the implementation of (i). The parameter α acts in the intensity domain similarly to d in the spatial domain: α controls the contribution of δI to $L(x)$, making the intensities of $\mathcal{B}(x)$ closer to $I(x)$ in the intensity domain more important than the others. The value $\alpha = 1$ is in this case admitted and used to weight equally the intensity variations at a given distance, regardless of their amount, since $c_0(\delta I(x,b)) = 1$ for any value of $\delta(x,b)$. When $\alpha < 1$, the relevance of $\delta I(x,b)$ decreases proportionally to the value $\delta I(x,b)$. The parameter b controls the value of f as d and δI grow away: the lower b, the lower the effects of high values of d and δI are on $L(x)$. Values of b greater than a and α make f increasing when d and δI grow and this trend has to be prevented because in contrast with both the requirements (i) and (ii). Therefore, b must be smaller than $\min\{a, \alpha\}$. Finally, we note that for some choices of α, a, b, the values of C may become smaller than zero: these values must be excluded from the computation of f, therefore C is cast to zero as described by Equation (1). This operation extends the range of variability of the parameters α, a, b defined by the inequalities (6), cutting down values of f for which the requirements (i) and (ii) are not fulfilled.

Examples of f obtained for different values of α, a, b are shown in Fig. 3 along with the corresponding enhancement by SuPeR-B. For all the parameter values used here, the yellowish color due to the light has been removed. In particular, for $\alpha = 0, -1$, the input image has been remarkably improved. A more deep analysis of the enhancement results of SuPeR-B is provided in the next Section.

3.3 Experiments

The performance of SuPeR-B has been measured on a database of 60 images with dark and bright regions at different proportions. These images have been taken from Pixabay (https://pixabay.com/), from private collections of the author and from databases published on the net and used to test and illustrate enhancement methods, e.g. (Wang et al., 2016), (Fu et al., 2016), (Li and Wu, 2018), (Ramirez Rivera et al., 2012). Some examples of such images are shown in Figures 3 and 4.

The results of SuPeR-B have been compared with the four enhancers briefly described in Section 2, i.e. Channel Division (Ramirez Rivera et al., 2012)[1], Fusion (Wang et al., 2016)[2], Backlit (Li and Wu, 2018)[3] along with an implementation of the multiscale Retinex implementation (Petro et al., 2014)[4].

The evaluation has been carried out by comparing three objective measures of image quality before and after applying the enhancement procedures. These measures, that are widely used to judge the perfor-

[1]Code: http://vision.khu.ac.kr/?page_id=551
[2]Code: https://xueyangfu.github.io/
[3]Code: https://github.com/7thChord/backlit
[4]Code: https://it.mathworks.com/matlabcentral/fileexchange/71386-multiscale-retinex

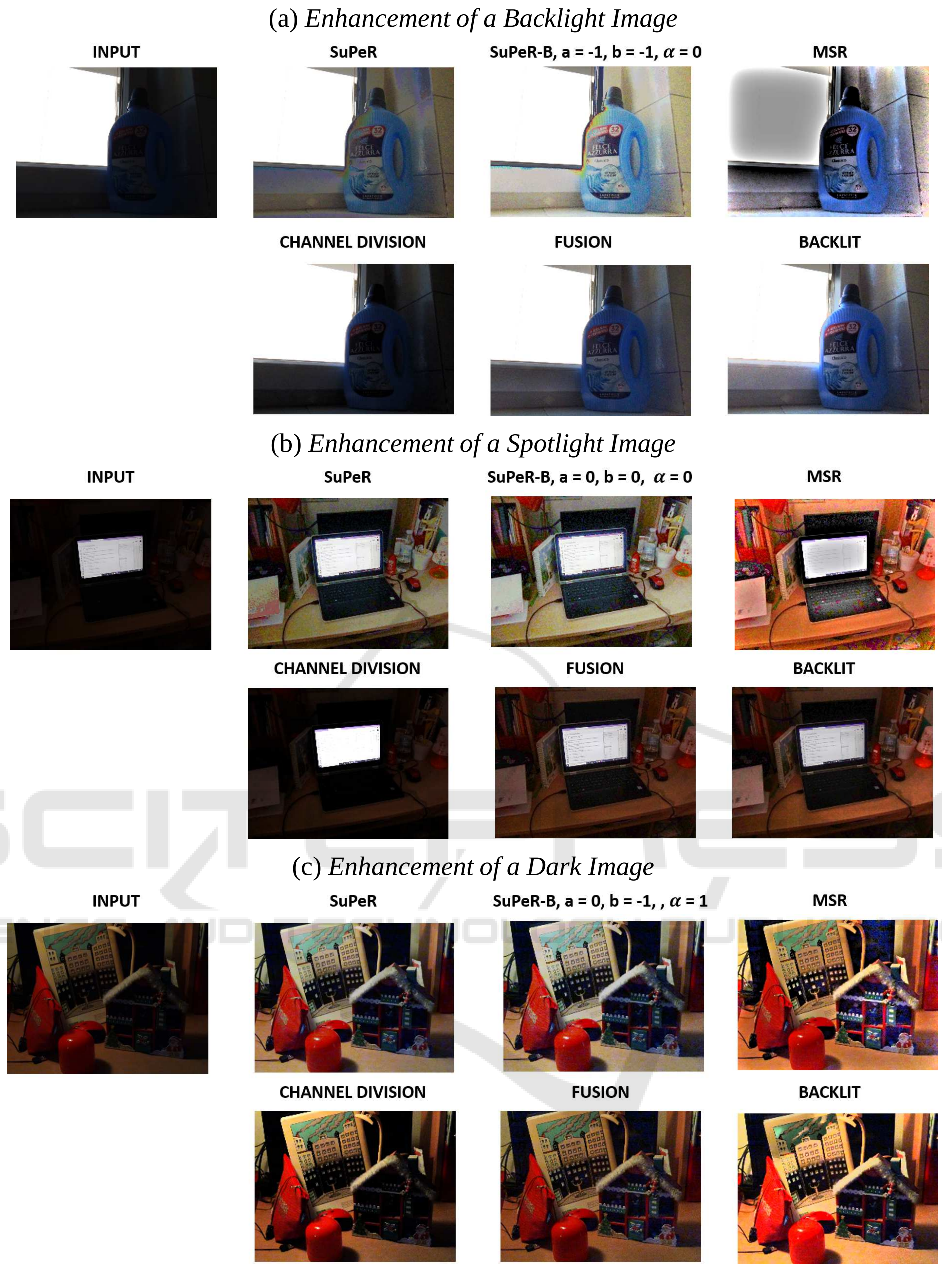

Figure 5: Comparison of enhancement results on a (a) backlight, (b) spotlight, (c) dark image.

mance of an enhancer, are computed on the brightness Br of any color image $\bar{I}$. Br is the gray-level image defined on S and obtained by averaging pixel-wise the channel intensities of $\bar{I}$, i.e.:

$$Br(x) = \frac{1}{3}\sum_{i=0}^{3} I_i(x) \qquad \forall\, x \in S, \tag{7}$$

where the I_is denote the color channels of $\bar{I}$.

The evaluation measures are defined as follows:

1. Flatness of the luminance probability density (f_0), i.e. the L^1 distance between the probability density function of Br and the uniform probability density function. The lower f_0, the lower the entropy of Br is. An image enhancer is expected to decrease the value of f_0 measured on the input image. In fact, this means that in the enhanced image the dynamic range of the colors is wider than in the input image, meaning that strong changes of light and brightness have been smoothed or even removed;
2. Mean value of the brightness (f_1), i.e. the average of the intensities of Br. Enhancement usually

increases the value of f_1, since it makes the input image brighter;

3. Mean value of the multi-resolution luminance contrast (f_2), i.e. a measure of the local variations of the Br values at multiple scales proposed in (Rizzi et al., 2004). f_2 is computed as follows. The image Br is sequentially half-scaled. For each rescaled version Br_s of Br, a pixel wise contrast $cBr_s(x)$ and a global contrast cBr_s are computed. Precisely, $cBr_s(x)$ is the average value of the absolute differences between the luminance $Br(x)$ at $x \in S$ and its 8-neighboring values, while cBR_s is the mean value of the $cBr_s(x)$'s. The measure f_2 is defined as the mean value of the cBr_s's over the number of scale factors. Any enhancer is expected to improve the visibility of the details, and thus to increase the value of f_2 with respect to that of the input image.

It is to note that the exact amount of the f_0, f_1 and f_2 depend on the image content. Images with already clear (unreadable, resp.) content have usually a low (high, resp.) value of f_0 and high (low, resp.) values of f_1 and f_2. What is important to evaluate an enhancer is the variation of the f_is after the enhancement: the less readable the image, the more relevant this variation is. Therefore, in this work the evaluation of the performances of the enhancers has been conducted by comparing the values of the f_is on the input and on the enhanced images. In addition, in order to provide more detailed information about the enhancement of the dark and bright regions, each input image has been segmented in two parts P_B and P_D. The f_is are thus computed on P_D and P_B instead of the whole image and indicated respectively with f_i^D and f_i^B. The segmentation has been performed by a thresholding procedure such that:

$$P_B = \{x \in S : Br(x) > \tau\} \quad (8)$$

$$P_D = \{x \in S : Br(x) \leq \tau\} \quad (9)$$

and $\tau = (\max_{x \in S} Br(x) - \min_{x \in S} Br(x))/2$. Basically, this segmentation partitions the input image in two regions with different luminance, with P_D darker than P_B. Despite naive, this segmentation allows to separate the frontlight and backlight regions of the input images in a sufficiently reliable way.

The current implementation of SuPeR-B is written in C++. The parameter n has been set to 100. The experiments have been repeated for different triplets of (a, b, α) obtained by varying a, b in $\{-1, 0\}$ and α in {-1, 0, 1} within the inequalities in formula (6). For $\alpha = 1$ and $a = b = 0$, SuPeR-B behaves like SuPeR. The function f, that is computed in the global processing phase, is discretized and represented as a matrix with size 100×100 to speed up the computational time for the enhancement, that is on average less than 1 second (on a notebook with Intel CORE i5 and operating system Windows 10). In these experiments, no median filter has been applied.

Table 1 reports the mean values of the performance evaluation measures computed on the whole image (f_is), on the bright region P_B (f_i^Bs) and on the dark region (f_i^Ds). Precisely, these values have been averaged over the number of dataset images.

All the enhancers improve the image content, by increasing the values of f_1 and f_2, while decreasing that of f_0. MSR achieves the best results in term of f_0 and f_2, but as already observed in the paper MSR suffers for halos generation (see Fig. 5(a) for an example). Channel Division outputs the worst results in terms of f_0 and f_1, returning an image still dark and with a remarkable difference between dark and bright regions, as proved by the gap between f_1^Bs and f_1^Ds (see for instance, Fig. 5(b) and (c)). In particular, Channel Division yields the highest value of f_1^B and the lowest value of f_1^D. For $\alpha = -1$, SuPeR-B outputs the highest values of f_1 and very low values of f_0, while in general, for any choice of the parameters α, a, b, it obtains values of f_2 close to those of Channel Division, SuPeR, Fusion and Backlit.

A similar trend is observed for the dark regions P_D. Again, the readability of the P_Ds is improved by all the enhancers, MSR grants the lowest f_0^D and the highest f_2^D, Channel Division reports the lowest f_0 and SuPeR-B performs similar to SuPeR, Fusion, Backlit and Channel Division in terms of contrast while reports for $\alpha = -1$ very high values of f_1^D and very low values of f_0^D.

The results on the P_Bs could appear quite surprising: in fact, here, apart from Channel Division, all the enhancers increase the brightness, but decrease the contrast, while the histogram flatness remains high and in many cases exceeds that measured on the input images. This behaviour is justified as follows. On the bright regions, enhancement is usually not necessary since the content is already clear. Anyway, the brightness of the P_Bs, that is on average very high on the input, is further increased by the enhancers, although much less than that of the P_Ds. The point is that in the enhanced images, the brightness is very close to its maximum value (i.e. 255) and consequently its histogram is peaked on right and thus far from an uniform probability density function (i.e. f_0^D is high). Moreover, in the dataset considered here, the bright regions are often almost uniform or present few or slight edges. The main contribution to the contrast of the P_Bs comes from their boundaries with the dark regions P_Ds. Despite the enhancers preserve these boundaries, they diminish their magnitude be-

Table 1: Evaluation of SuPeR-B in comparison with other enhancers. For SuPeR-B the triplet (α, a, b) is reported.

Algorithm	f_0 $[\times 10^{-3}]$	f_1	f_2	f_0^B $[\times 10^{-3}]$	f_1^B	f_2^B	f_0^D $[\times 10^{-3}]$	f_1^D	f_2^D
INPUT	4.17	67.79	16.56	4.92	179.33	33.07	5.26	29.74	14.58
MSR	2.21	118.12	28.66	4.69	185.23	30.51	2.35	97.24	31.66
CHANNEL DIVISION	4.04	89.45	21.60	5.86	228.30	36.03	4.53	41.49	20.97
FUSION	3.23	89.54	20.52	4.84	189.51	32.27	4.22	56.82	19.82
BACKLIT	2.78	100.65	22.79	4.68	189.30	28.35	3.47	73.26	23.55
SuPeR	3.41	104.95	19.95	5.40	210.54	29.78	4.25	65.56	19.63
SuPeR-B (1, 0, -1)	3.34	106.64	20.06	5.40	210.78	29.61	4.17	67.71	19.31
SuPeR-B (1, -1, -1)	3.31	105.32	20.17	5.37	210.51	30.07	4.15	66.16	19.87
SuPeR-B (0, 0, 0)	3.17	119.31	20.05	5.51	213.03	26.53	3.95	84.12	20.46
SuPeR-B (0, 0, -1)	3.02	123.13	20.19	5.53	213.47	25.84	3.75	89.81	20.83
SuPeR-B (0, -1, -1)	2.89	123.48	20.48	5.50	213.43	26.14	3.61	90.02	21.15
SuPeR-B (-1, 0, -1)	2.92	141.15	20.08	5.88	218.83	20.42	3.26	113.97	22.17
SuPeR-B (-1, -1, -1)	2.78	142.35	20.33	5.87	219.29	20.28	3.05	115.44	22.55

cause they increase the brightness of P_D: this causes the decrease of f_2^D. Channel Division represents an exception, since the trend of the f_i^Bs is the same of those of the f_i^D and of the f_is. Nevertheless, as already mentioned above, among the enhancers, this algorithm achieves unsatisfactory results, with still dark images characterized by high gap between f_1^B and f_1^D. Regarding SuPeR-B, it is to note that the values of f_2^Bs it reports are in general smaller than those of the other enhancers. This is because, in line with Retinex theory and as well as SuPeR, SuPeR-B tends to remove slight edges, that are often present in the bright regions but that are considered to be irrelevant for understanding the main content of the scene. In particular, the low values of f_2^B indicate a scarce presence of edges on P_B, and this justifies the high values of f_0^B.

Figure 3 shows that different values of α, a, b provide very different enhancement results. In accordance with Equation (3), the parameters α and a tune respectively the contributions to $L_B(x)$ of the intensity differences and of the spatial differences, while the parameter b controls both these contributions when the intensity and spatial differences are large. In particular, lowering α increases the weight of the small intensity differences and this generally higher the values of f_1 and f_2, especially in dark areas. For any fixed value of α, low values of a increase the spatial locality of the algorithm, yielding again high values of f_1 and f_2. Low values of b further decreases the weight of high intensity and spatial differences. Therefore, the lower α, a, b, the higher the locality of the algorithm in spatial and intensity domains is.

According to such observations, when $\alpha = 1$, SuPeR-B behaves similarly to SuPeR: in fact, in this case the contribution to $L_B(x)$ of the intensity variations does not depend on their amount and the Coon patch is close to a stepway, whose slope is determined by the spatial terms. As already mentioned before, for $\alpha = 1, a = b = 0$, SuPeR-B implements SuPeR, where the locality of the image processing depends exclusively on the spatial features. Figure 3(a) show an example of image enhancement obtained by tuning only the contribution of the spatial differences with $\alpha = 1, a = b = -20$: the details of the giraffe become a little more visible, but the giraffe is still quite dark and its borders are very tick, so that the results is unsatisfactory. Better results are obtained by tuning the algorithm locality also in the intensity domain, as shown by the enhanced pictures of Figure 3(a) obtained for $\alpha = -1, -2$.

In general, the experiments show that for backlight and spotlight images the best performance is achieved by tuning the differences in both the spatial and intensity domains.

Finally, differently from the algorithms considered here exept for SuPeR, SuPeR-B exhibits the important property of lowering the color cast due to the illumination, returning an enhanced image robust to changes of light. This property is clearly illustrated in Figure 5(c) and it is of great importance for many machine vision algorithms, such as illumination invariant feature matching and people/object recognition in real-world scenarios.

4 CONCLUSIONS

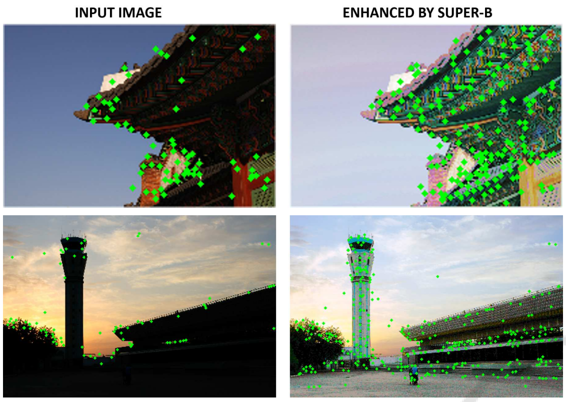

Figure 6: Examples of usage of SuPeR-B as a pre-processing step of the key-point detector in (Lowe, 2004): the enhancement of the input images (on left) grants a better detection of key-points (in green color, on right) in the dark regions.

The Retinex inspired spatial color algorithm SuPeR-B proposes a novel and efficacious technique to enhance images captured under difficult light, in particular under backlight and local, not diffused spotlights that hamper understanding the image content. SuPeR-B basically implements a bilateral processing of the channel intensities of the image pixels that enable brightening dark regions, smoothing color casts due to the light, while preserving important edges. The bilateral processing is modeled by the weighting function f, that here has been expressed as a Coon surface bounded by lines. The experiments proved that this choice of f provides a satisfactory level of enhancement, also in comparison with other algorithms at the state-of-the-art. The input images are remarkably improved by SuPeR-B, which increase their brightness and detail visibility, especially in the dark areas, while decrease the overall color distribution entropy.

As mentioned in Section 3, there are many expressions for f: determining the equation of f and the values of its parameters most suitable for enhancing an image within a given task is a critical point for a reliable and aware usage of SuPeR. In general, in the current implementation of SuPeR-B, the choice of the values of α, a, b to be input to SuPeR-B should be guided by the applications at the hand as well as by the image content. For instance, for human inspection of the content of the giraffe image in Figure 3, the values (α, a, b) = (0, 0, 0) and (1, 0, -1) perform poorly in comparison with the others. In other cases, like for (α, a, b) = (-1, -1, -1), the brightness of the enhanced image is too high and the image content is washed out or over-enhanced. As a conclusion, the set-up of f is an open issue to be investigated in the future.

Future work will also include the usage of SuPeR-B within computer vision tasks that require to process high quality images, as for instance image description and matching. In this respect, as pointed out in (Lecca et al., 2019) and as illustrated by the examples in Figure 6, the improvement of visual characteristics such as brightness, contrast and color distribution does not only grant a better visibility and readability of the image content for humans, but it is also fundamental to improve the detection of image key-points (here performed by SIFT (Lowe, 2004)) and thus to increase the performance of algorithms for description and matching of images regardless of their illumination conditions.

REFERENCES

Banić, N. and Lončarić, S. (2013). Light Random Sprays Retinex: exploiting the noisy illumination estimation. *IEEE Signal Processing Letters*, 20(12):1240–1243.

Banić, N. and Lončarić, S. (2015). Firefly: A hardware-friendly real-time local brightness adjustment method. In *2015 IEEE International Conference on Image Processing (ICIP)*, pages 3951–3955.

Finlayson, G. D., Drew, M. S., and Funt, B. V. (1994). Color constancy: generalized diagonal transforms suffice. *JOSA A*, 11(11):3011–3019.

Fu, X., Zeng, D., Huang, Y., Liao, Y., Ding, X., and Paisley, J. (2016). A fusion-based enhancing method for weakly illuminated images. *Signal Processing*, 129:82–96.

Gianini, G., Manenti, A., and Rizzi, A. (2014). QBRIX: a quantile-based approach to Retinex. *J. Opt. Soc. Am. A*, 31(12):2663–2673.

Guo, X., Li, Y., and Ling, H. (2017). Lime: Low-light image enhancement via illumination map estimation. *IEEE Transactions on Image Processing*, 26(2):982–993.

Jiang, Y., Gong, X., Liu, D., Cheng, Y., Fang, C., Shen, X., Yang, J., Zhou, P., and Wang, Z. (2019). Enlightengan: Deep light enhancement without paired supervision. *arXiv preprint arXiv:1906.06972*.

Jobson, D. J., Rahman, Z., and Woodell, G. A. (1997). Properties and performance of a center/surround retinex. *IEEE Transactions on Image Processing*, 6(3):451–462.

Jobson, D. J., Rahman, Z.-u., and Woodell, G. A. (1997). 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.

Lecca, M. (2014). *On the von Kries Model: Estimation, Dependence on Light and Device, and Applications*, pages 95–135. Springer Netherlands, Dordrecht.

Lecca, M. (2018). STAR: A segmentation-based approximation of point-based sampling Milano Retinex for color image enhancement. *IEEE Transactions on Image Processing*, 27(12):5802–5812.

Lecca, M. and Messelodi, S. (2019). SuPeR: Milano Retinex implementation exploiting a regular image grid. *J. Opt. Soc. Am. A*, 36(8):1423–1432.

Lecca, M., Simone, G., Bonanomi, C., and Rizzi, A. (2018). Point-based spatial colour sampling in milano-retinex: a survey. *IET Image Processing*, 12(6):833–849.

Lecca, M., Torresani, A., and Remondino, F. (2019). On image enhancement for unsupervised image description and matching. In Ricci, E., Rota Bulò, S., Snoek, C., Lanz, O., Messelodi, S., and Sebe, N., editors, *Image Analysis and Processing – ICIAP 2019*, pages 82–92, Cham. Springer International Publishing.

Li, S., Cheng, Q. S., and Zhang, J. (2020). Deep multi-path low-light image enhancement. In *2020 IEEE Conference on Multimedia Information Processing and Retrieval (MIPR)*, pages 91–96.

Li, Z. and Wu, X. (2018). Learning-based restoration of backlit images. *IEEE Transactions on Image Processing*, 27(2):976–986.

Lowe, D. G. (2004). Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60(2):91–110.

Lv, F., Lu, F., Wu, J., and Lim, C. (2018). MBLLEN: Low-light image/video enhancement using cnns. In *BMVC*, page 220.

M. Lecca, A. Rizzi, and R.P. Serapioni (2017). GREAT: a gradient-based color-sampling scheme for Retinex. *J. Opt. Soc. Am. A*, 34(4):513–522.

Morel, J. M., Petro, A. B., and Sbert, C. (2010). A PDE formalization of Retinex theory. *IEEE Transactions on Image Processing*, 19(11):2825–2837.

Petro, A. B., Sbert, C., and Morel, J.-M. (2014). Multiscale retinex. *Image Processing On Line*, pages 71–88.

Ramirez Rivera, A., Byungyong Ryu, and Chae, O. (2012). Content-aware dark image enhancement through channel division. *IEEE Transactions on Image Processing*, 21(9):3967–3980.

Rizzi, A., Algeri, T., Medeghini, G., and Marini, D. (2004). A proposal for contrast measure in digital images. In *CGIV 2004 - 2nd European Conference on Color in Graphics, Imaging, and Vision and 6th Int. Symposium on Multispectral Color Science*, pages 187–192, Aachen.

Rizzi, A. and Bonanomi, C. (2017). Milano retinex family. *Journal of Electronic Imaging*, 26(3):031207.

Tsai, C.-M. and Yeh, Z.-M. (2010). Contrast compensation by fuzzy classification and image illumination analysis for back-lit and front-lit color face images. *IEEE Transactions on Consumer Electronics*, 56(3):1570–1578.

Wang, Q., Fu, X., Zhang, X., and Ding, X. (2016). A fusion-based method for single backlit image enhancement. In *2016 IEEE International Conference on Image Processing (ICIP)*, pages 4077–4081.

Wang, S., Zheng, J., Hu, H.-M., and Li, B. (2013). Naturalness preserved enhancement algorithm for non-uniform illumination images. *IEEE Transactions on Image Processing*, 22(9):3538–3548.

Wei, C., Wang, W., Yang, W., and Liu, J. (2018). Deep retinex decomposition for low-light enhancement. In *BMVC*.

Zuiderveld, K. (1994). Contrast limited adaptive histogram equalization. In *Graphics gems IV*, pages 474–485. Academic Press Professional, Inc.