

# Gradient attenuation as an emergent property of reset-based Retinex models

Gabriele Gianini

EBTIC/Khalifa University of Science and Technology  
Abu Dhabi, UAE  
Università degli Studi di Milano  
Milano, Italy  
gabriele.gianini@unimi.it

Leopold Ghemmogne Fossi

LIRIS/DRIM INSA of Lyon  
Villeurbanne, France  
Università degli Studi di Milano  
Milano, Italy  
leopold.fossi@unimi.it

Corrado Mio

Dipartimento di Informatica  
"Giovanni degli Antoni"  
Università degli Studi di Milano  
Milano, Italy  
corrado.mio@unimi.it

Alessandro Rizzi

Dipartimento di Informatica  
"Giovanni degli Antoni"  
Università degli Studi di Milano  
Milano, Italy  
alessandro.rizzi@unimi.it

## ABSTRACT

The Retinex image filtering algorithms have been inspired by experimental findings on the behavior of the Human Vision System. They are known to locally adjust image color and contrast by preserving edges and attenuating gradients. In a reference formulation of the algorithm by Land and McCann, edge preservation and gradient attenuation are granted by two ad-hoc mechanisms: called respectively *reset* (the distinctive feature of all the Retinex algorithms) and *thresholding*. A somehow unanticipated finding is that gradient attenuation is also observed with algorithm variants that do not include the latter mechanism, which was explicitly devised to implement gradient attenuation. In this work, we provide an analytic demonstration of the capability of Retinex models to attenuate gradients using only the "reset" mechanism, combined with the local character of the mutual pixel influences. We show that this capability is an emergent property of all the reset-based Retinex models.

## CCS CONCEPTS

• **Applied computing** → Imaging.

## KEYWORDS

Retinex models; Reset mechanism; Thresholding mechanism; Gradient attenuation.

## ACM Reference Format:

Gabriele Gianini, Corrado Mio, Leopold Ghemmogne Fossi, and Alessandro Rizzi. 2019. Gradient attenuation as an emergent property of reset-based Retinex models. In *MEDES '19: 11th International Conference on Management of Digital EcoSystems (MEDES '19)*, November 12–14, 2019, Limassol, Cyprus. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3297662.3365784>

## 1 INTRODUCTION

The Spatial Color Algorithms (SCA) [26] family, for which Retinex is the founding model, is a group of algorithms inspired by the Human Vision System (HVS), built on the principle – suggested by experimental observations – that the color sensation at a point is not determined simply by the color stimulus at that point, but rather by the relationship between that stimulus and a wider set of stimuli, ranging from the surrounding area to the rest of the image. The algorithms of this family differ from one another mainly in the mechanisms by which the information from the surrounding is selected, and aggregated to yield the output color.

The Retinex algorithm in [10] – the earliest model of the family – has originated many close variants [2, 3, 8, 9, 11, 13–15, 17, 18, 20, 21, 23, 24, 28–30, 32], some of which focus on the way to select and use the neighboring pixels [4–7, 12, 31].

In the original Retinex [20] and the Milano Retinex family [22, 28] (a detailed description of their differences can be found in [19, 27]), the information from the surrounding is collected by statistical sampling – the samples are typically paths [14–17] or sprays [1, 23]) – then sample-wise compared to the center/target pixel using an operator called *reset*, which in the simplest cases corresponds to the "max" operator (more details in the following sections). The reset outputs a value, which represents the contribution of that particular sample.

The output at a target pixel is obtained (in each chromatic channel) by re-scaling the input with respect to a *reference white*, computed as the harmonic average of the different samples' contributions. On the one hand, the detailed definition of the sampling mechanism controls the intensity of the influence of the pixels at different distances from the target pixel and determines how much

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

MEDES '19, November 12–14, 2019, Limassol, Cyprus

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6238-2/19/11...\$15.00

<https://doi.org/10.1145/3297662.3365784>

*local* the resulting filtering will be, on the other hand, the *reset* grants edge preservation.

Along with edge preservation, also gradient attenuation had been observed in perceptual experiments [17]: this observation prompted Land and McCann, in their seminal paper [10], to introduce an additional ad-hoc mechanism for smooth gradient suppression – called *thresholding*. The thresholding, was also meant to render the finite sensitivity of the HVS to intensity differences of contiguous pixels, which was regarded as the main cause of the perceived gradient attenuation.

To some surprise, even turning off the thresholding – i.e. using only the reset mechanism – gradients turned out to be attenuated. In fact, this effect can be appreciated in all the versions of the reset-based Retinex algorithms, even in those that, by construction, cannot embody any form of thresholding (such as RSR [23]). This finding was unanticipated, since the core mechanism – the *reset* – was devised to preserve edges, at the same time rescaling w.r.t. local maxima, but was not built for gradient attenuation.

The present work aims at bringing some clarification to this puzzle. We demonstrate analytically that the *reset* mechanism, combined with the *locality* of the pixel influence does, indeed, cause gradient attenuation. In other words, we prove that gradient attenuation is an emergent property of reset-based Retinex models (the emergent character of the effect is discussed at the end of the paper).

To this purpose, we examine a sequence of Spatial Color Algorithms of increasing complexity: first we recall the definition of the main Retinex mechanisms and show how they are realized in the different Retinex variants (Section 2), then (Section 3) we focus on gradient attenuation: first, using a stylized algorithm (Local Scale-to-Max), we demonstrate analytically how *reset* and *locality*, combined, yield gradient attenuation (Subsection 3.1) and indicate how the same combination of mechanisms can explain gradient attenuation in all the Retinex algorithms based only on the reset (i.e. without thresholding) (Subsection 3.2). A discussion of gradient attenuation as an emergent property of the model (Section 4) concludes the paper.

## 2 RETINEX AND THE SPATIAL COLOR ALGORITHMS

### 2.1 The defining mechanisms of retinex

To recall the defining Retinex mechanisms, we use the path-based Retinex. Consider a path  $\gamma$  taken from the collection of all the possible oriented paths to  $x$ : it defines an ordered list of pixels from the image. Let us denote by  $t$  the step count and indicate by  $r_t$  the ratio of the intensities  $i(x_t)$  and  $i(x_{t-1})$  of two consecutive pixels along the path: for  $t > 0$   $r_t \equiv \frac{i(x_t)}{i(x_{t-1})}$ . By convention  $r_0 \equiv 1$ . If the pixel at step  $t$  is darker than the one at  $(t - 1)$ , then  $r_0 < 1$ .

**2.1.1 Thresholding.** The *threshold* mechanism is introduced by means of a *threshold function*  $\theta^{(\epsilon)}(\cdot)$ , defined for  $s \in \mathbb{R}$  by

$$\theta^{(\epsilon)}(s) \equiv \begin{cases} 1 & \text{if } (1 - \epsilon) < s < (1 + \epsilon) \\ s & \text{otherwise} \end{cases} \quad (1)$$

where  $\epsilon > 0$  is a parameter of the model. This operator filters out smooth intensity changes, but preserves sharper intensity changes, e.g. edges.

**2.1.2 Reset.** The reset mechanism is introduced through a "cap" function  $\rho(\cdot)$

$$\rho(s) \equiv \begin{cases} s & \text{if } s < 1 \\ 1 & \text{otherwise} \end{cases} \quad (2)$$

The contribution of each path is computed by means of a composite function of the path-ordered intensities, which, at every step, applies the threshold function to the current intensity ratio and then the reset filter to the chain product of ratios; such a function is defined *recursively* as follows:

$$C^{(\epsilon)}(t) \equiv \begin{cases} 1 & \text{for } t = 0 \\ \rho \left( \theta^{(\epsilon)}(r_t) \times C(t-1) \right) & \text{for } 1 \leq t \leq n \end{cases}$$

The contribution by a path to the target pixel  $x$  is defined as the value at the last step  $C^{(\epsilon)}(n)$  (at which it gets to  $x$ ). The output  $o(x)$  at a pixel  $x$  is computed, in Milano Retinex, by averaging those contributions over a large set of paths.

It can be easily shown that switching off the thresholding mechanism (setting  $\epsilon = 0$ ) is equivalent to computing the maximum of a path and to compute the output as the input times the harmonic average of the collection of maxima. Most variants of Retinex drop the thresholding and use only the reset, thus, in fact, compute the output by harmonic averaging the maxima of their samples.

### 2.2 Stylized algorithms: Scale-to-Max and Local Scale-to-Max

Hereafter we represent the input gray-scale image by the function  $i : x \in X \rightarrow i(x)$ , defined on the support  $X$ , and the output image by  $o : x \in X \rightarrow o(x)$ .

**2.2.1 Scale-to-Max.** The algorithm Scale-to-Max is a simple algorithm, which rescales the input using, as a single reference white, the brightness of the brightest pixel from those forming the image; it is also known as von Kries filter [25, 33]. In Scale-to-Max the filtered output image  $o$  at any point  $x_0$  is computed as

$$o(x_0) = \frac{i(x_0)}{\max\{i(x)\}_{x \in X}} = \frac{i(x_0)}{i_X^{max}}$$

Since  $X$  is the whole image support, the algorithm uses a global filtering criterion.

**2.2.2 Local Scale-to-Max.** The algorithm Local Scale-to-Max introduces some locality in the correction: the input is rescaled using as reference white the value of the brightest pixel in a circular neighborhood, for a predefined radius  $\Delta$ . The filtered output image  $o$  at any point  $x_0$  is computed as

$$o(x_0) = \frac{i(x_0)}{\max\{i(x)\}_{x \in \Omega_\Delta(x_0)}} = \frac{i(x_0)}{i_\Omega^{max}}$$

with  $\Omega_\Delta(x_0) = \{x \in X : |x - x_0| < \Delta\}$ , for a fixed radius  $\Delta > 0$  shorter than the image diagonal. Clearly, this algorithm uses a local filtering criterion.

This is considered a stylized or toy algorithm because it does not provide a satisfactory filtering: the sharp boundary between the

neighboring region  $\Omega$  and its complement can create, in the vicinity of the brightest areas, abrupt changes in the reference white and determine the formation of artifacts. However, this algorithm comprises all the qualitative traits necessary for deriving the analytical form of the gradient attenuation, as we are going to show in Section 3.

Notice that both Scale-to-Max and Local Scale-to-Max use all the pixel of a given area to find the reference white: the former considers all the pixels of the whole image support  $X$ , while the latter considers all the pixels of the neighborhood region  $\Omega_\Delta(x_0)$  of the pixel at  $x_0$ . The algorithms described below, on the contrary use a subset of pixels from  $X$  chosen by statistical sampling.

### 2.3 Sampling Retinex algorithms using only the reset mechanisms

Most Retinex algorithms use samples to explore the neighborhood of a point: those point sets are obtained typically by random paths or by generating isotropic samples of points (the so-called *sprays*) around the target pixel. The typical structure of the sampling-based algorithms of the Milano Retinex family, for computing the output  $o(x_0)$  at a target point  $x_0$  is the following. Generate a sample of points near the target; collect their intensities into a set  $S$ ; add the pixel intensity  $i(x_0)$  to the set, obtaining an augmented sample  $S^*$ . At this point the algorithms *that do not implement the thresholding* find the maximum intensity of  $S^*$ , i.e.  $i_S^{max} \equiv \max\{i(x_0), \max\{i \in S\}\}$  and use it as the contribution by  $S$ ; the operation is repeated for many samples  $S$  and eventually the harmonic average is taken as reference white. Overall

$$o(x_0) = i(x_0) \left\langle \frac{1}{i_S^{max}} \right\rangle$$

where  $\langle \cdot \rangle$  denotes the average over the samples, i.e. for a number  $N$  of samples

$$\left\langle \frac{1}{i_S^{max}} \right\rangle = \frac{1}{N} \sum_S \frac{1}{i_S^{max}}$$

The number  $N$  of samples is determined so as to provide a reasonable exploration of all the starting regions in the image and to grant that the final correction is affected as little as possible by statistical noise.<sup>1</sup>

Notice, for future reference, that this statistical process can be interpreted as the computation of the sampling average of the reciprocal of the maximum.

**2.3.1 Path-based Milano Retinex.** In its standard definition [22], the path-based Milano Retinex collects the samples in the following way. One generates starting points uniformly at random over the image, then generates a memoryless random walk  $\gamma$  over the set of pixels until it meets the target  $x_0$ : at that point, it stops, and the maximum intensity is used as  $i_\gamma^{max}$  for that sample. Many path-based variants are based on this one [13].

<sup>1</sup>Statistical noise becomes perceivable when nearby pixels with similar intensities  $i'$  and  $i''$  yield remarkably different outputs  $o'$  and  $o''$ , due to the difference in the sample collection: the set of maximum intensities found by the collection of  $N$  samples for  $i'$  is so different from the set maximum intensities found by the collection of  $N$  samples for  $i''$ , that the harmonic averages are sensitively different.

**2.3.2 Random Spray Retinex (RSR).** To reduce the computational effort and the redundancy in the samples collected using paths, a version of Milano Retinex called Random Spray Retinex or RSR has been devised [23]. In this algorithm, each sample consists of a set of points collected randomly around the target  $x_0$  according to a decreasing sampling density. The set of points is called *spray*. The number  $n$  of points in the spray and the number  $N$  of sprays are parameters that can be tuned to reduce as much as possible statistical noise. A version of this algorithm based on fuzzy set theory was also created to further reduce the computational cost [7].

### 2.4 Algorithms based on probabilistic models of sampling

To remove completely the statistical noise generated by the sampling based algorithms, some algorithms have been devised, which modeled the sampling as a probabilistic process [4–6, 8]: in those algorithms the averages are not estimated statistically, but computed exactly, based on a probabilistic model.

**2.4.1 QBRIX.** The simplest Retinex probabilistic modeling algorithm is QBRIX, short for Quantile-Based Retinex [6]. The starting observation for the construction of QBRIX is that when computing the sampling average of the reciprocal of the max of a distribution one obtains a high order quantile of the distribution. Consequently, in place of proceeding by sampling, one can resort to the construction of the intensity histogram and proceed directly to the computation of a high-order quartile. As the parameters in other algorithms (such as the number  $n$  of points in a spray for RSR) the quantile order  $q$  can be tuned based on the image characteristics. The intensity histogram can be built using all the pixels from the whole image, with the same weight (Global QBRIX), or changing the weights of each point falling into the histogram, based on its distance from the target pixel, as in RSR (Local QBRIX).

The prescription for the global version of QBRIX is the following. Create an histogram  $f(i)$  containing the intensities  $i(x)$  for  $x \in X$ ; consider its cumulative  $F(i)$ ; compute the (pseudo)-inverse of the cumulative i.e. the quantile function<sup>2</sup>  $i_q \equiv F^{(-1)}(q) = \{i : F(i) = q\}$ . Once found  $i_q$ , if  $i(x_0) < i_q$  set

$$o(x_0) = \frac{i(x_0)}{i_q}$$

Set  $o(x_0) = 1$  elsewhere (set to white every value above the reference quantile).

**2.4.2 Local QBRIX.** In the local version, the histogram is built with reference to the target point  $x_0$ , using a weight function  $\lambda(|x - x_0|)$  decreasing with the distance, such as  $h(|x - x_0|) \propto (x - x_0)^{-2}$ , to weight each histogram entry. Thus the reference quantile for a pixel strongly depends on the closer neighborhood. Denoting by  $i_q^\lambda(x_0)$  such local quantile reference one can compute the output as

$$o(x_0) = \frac{i(x_0)}{i_q^\lambda(x_0)}$$

if  $i(x_0) < i_q^\lambda(x_0)$ , and set  $i(x_0) = 1$  elsewhere.

<sup>2</sup>The pre-image of  $q$  can be a set of intensities, however, using in advance a suitable regularization of  $F$  one can get always a single intensity as  $i_q$ .

**2.4.3 RSR-P and ReMark.** The full-fledged probabilistic version of RSR is RSR-P [5]: it models, step by step, the probabilistic processes at work in RSR and it computes the sampling average the whole set of pixels.

Similarly, ReMark [8] represents the probabilistic version of the path-based Milano Retinex. The model relies on the computation of probability fluxes with suitable Absorbing Markov Chains. The computations necessary to generate the output are rather involved, requiring the repeated solution of large linear systems.

Although the two models do not allow in general closed-form analytic computations, they are often reasonably well approximated by QBRIX: in some simple cases, both RSR-P and ReMark coincide with Local QBRIX, for a suitable value of  $q$  [5]. This is precisely the case of the gradient images used in the following section.

### 3 GRADIENT ATTENUATION

Here we model analytically the mechanism of gradient attenuation that is originated by the reset mechanism in various algorithms.

#### 3.1 Effect of Local Scale-to-Max on a gradient

In the local version of Scale-to-Max one can see two stylized versions of the mechanisms that are at work in the reset-based Retinex models:

- 1) a rescaling qualitatively similar to the one produced by the reset mechanism (based on the max over a full area, instead of the max over a sample) and
- 2) the locality of the correction algorithm (the max is computed for an area within a given radius from the target pixel, rather than by using all the pixels and a distance-dependent weight).

Consider a simple input image  $i(x, y) = i(x)$  consisting in a ramp of slope  $\alpha$ , over an interval  $[x_{min}, x_{max}]$  of the spatial variable  $x$  (we disregard, here the spatial variable  $y$ , that has no influence on the output):

$$i(x) = c + \alpha x$$

where  $c$  can be any non-negative constant. Let the range  $\Delta$ , specifying the radius of the neighborhood  $\Omega$ , be given. The maximum in the neighborhood of a target point  $x_0$  is – due to monotonicity of the ramp – at  $x_0 + \Delta$  and has value  $i_{\Omega}^{max} = i(x_0 + \Delta) = i(x_0) + \alpha\Delta$  (provided that  $x_0$  is far enough from the border). Thus the output at  $x$  for this special image (provided that  $x \leq (x_{max} - \Delta)$ ) is

$$o(x) = \frac{i(x)}{i(x) + \alpha\Delta} = \frac{1}{1 + \frac{\alpha\Delta}{i(x)}} = \frac{1}{1 + \frac{\Delta}{(c/\alpha)+x}}$$

The key observation of the present discussion is the following: for  $(c/\alpha)+x$  high w.r.t. to  $\Delta$ , this function tends to a *horizontal line*. The mechanism is described with an even simpler function in Figure 1, while the effect on a gradient image is shown in Figure 2.

#### 3.2 Effect of other Retinex algorithms on a gradient image

We discuss the gradient reduction capability of the other algorithms – introduced in the previous section – without using analytical tools: the analytic treatment of their effect on a simple slope function is possible but rather involved due to unavoidable border

effects. Indeed the weight functions involved in RSR, RSR-P, and QBRIX have radial symmetry, while the input image is rectangular. However, the qualitative behavior of those algorithms can be inferred by the salient characteristics of their definitions.

**3.2.1 Local QBRIX.** The effect of Local QBRIX on a gradient is very similar to the one of Local Scale-to-Max. This is a consequence of the regularity of the slope input function  $i(x) = c + \alpha x$ .

Computing the output of Local QBRIX amounts to building a pixel-specific histogram of the weighted intensities in the image – the weight function being a decreasing function  $h$  of the distance from the target pixel – then picking the  $q$ -quantile (of a high order  $q$ ) of such an histogram: the intensity  $i_q^\lambda(x_0)$ , to compute the output  $o(x_0) = i(x_0)/i_q^\lambda(x_0)$ .

Thanks to the regularity of the input image  $i(x)$  – given a  $x_0$  sufficiently far from the border, the translation of  $x_0$  by a small displacement  $\delta$  toward the right causes the translation of the  $i_q^\lambda(x_0)$  by an amount approximately equal to  $\alpha\delta$ . Despite this approximation, the reference white  $i_q^\lambda(x_0)$  translates almost rigidly with the reference point intensity: an effect that in Local Scale-to-Max was able to produce gradient attenuation. Thus, also Local QBRIX causes gradient attenuation.

**3.2.2 RSR-P and RSR.** By construction, Local QBRIX is an approximation of RSR-P [5], which in turn is the probabilistic model of the process enacted by RSR. The main difference lies in the fact that QBRIX uses a fixed quantile order  $q$  for all the image, whereas the intensity identified by RSR-P as a reference white corresponds in general to a different quantile for every pixel.

In the case of a slope function however the relative distribution of the intensities with respect to the target pixel intensity is (disregarding border effects) homogeneous over the image: in other words, all the pixel-specific histograms have the same shape (even if they are translated w.r.t. one another). So, every reference identified by RSR-P corresponds to the same quantile value  $q$  for the pixel-specific histogram. In short: for the gradient image, Local QBRIX coincides with RSR-P. It follows that RSR-P causes naturally gradient attenuation, and that so does RSR.

**3.2.3 ReMark and Path-based Milano Retinex.** RSR was built as an approximation of the Path-based Milano Retinex (version without thresholding). It follows that the respective probabilistic versions are approximations of one another, i.e. RSR-P is an approximation of ReMark.

The points forming the sprays were intended as an approximation of the sample of points visited by a path (typically a memory-less random walk): the spray is less redundant, thus more computationally efficient.

The main qualitative difference between the sprays and the paths, in terms of information carried to the target pixel, is the following. A path, in order to go from one point to the target pixel, has to visit a chain of pixels crossing the region in between, so every information from a far region is necessarily mediated by information of intermediate regions. If there is a tiny bright region close to the target – let us call this region a geometric "barrier" – by crossing it, the path will discard the information it carries from farther regions, due to the reset mechanism (it will start carrying the new maximum brightness of this "barrier"). On the contrary, the sprays




Figure 1: Illustration of the effect of Scale-to-Max on a slope with constant gradient  $\alpha$ .




Figure 2: Scan-lines and images illustrating the effect of Scale-to-Max on a slope with constant gradient  $\alpha$ . Left: at the bottom, the input (blue line), just above, the reference white level (green line), at the top, the output (red line). We used an image rescaled in the interval  $i \in [0, 255]$  and  $[x_{min}, x_{max}] = [0, 100]$  The expression for the output holds up to  $x \leq x_{max} - \Delta$ , then border effects set up. Right: at the bottom, the input image, at the top, the output image. One can notice that, except for the region with border effects, the slope of the output is considerably reduced w.r.t. the one in the input.

sample all the points of the surrounding and a bright region has a probability of being sampled and becoming the maximum only proportional to its extension: sprays can "fly over barriers". This is just an example: in general, the paths are more sensitive to the asperities of the brightness landscape.

In the case of the gradient image, however, there are no barriers, and no asperities at all, therefore the qualitative behavior of spray sampling and memoryless random walk sampling are very similar. Other minor qualitative differences, between paths and sprays, materialize only on landscapes endowed at least by some minimal complexity, while they disappear if the image is defined by a linearly increasing function. In short, on the gradient image, RSR and RSR-P are good approximations of Path-based Milano Retinex (without thresholding) and ReMark, respectively. It follows that also the latter two algorithms are endowed with gradient attenuation capabilities.

#### 4 DISCUSSION AND CONCLUSIONS

We have shown that gradient attenuation is an intrinsic property of the reset-based Retinex algorithms, that arises even in the absence of the ad-hoc thresholding mechanism. This turns out to be an emergent property of the Retinex algorithms. Here we use the

term "microscopic" for the path-level (or spray level) description of the mechanisms, and the term "macroscopic" for the whole-image level description of the effects. This allows us to speak of *reductionist properties* to indicate those present at a microscopic level and preserved at the macroscopic level and to speak of *emergent properties* to indicate those that are properties of the operation of the algorithm at the macroscopic level, but are not of the algorithm at the microscopic level. It is apparent that while edge preservation can be obtained by a reductionist approach by means of the reset mechanism, gradient suppression is an emergent property of the algorithm. The path-wise reset mechanism results in the discovery of "jumps" in the reference white at "microscopic" level, that are preserved by the subsequent averaging at the macroscopic level. On the other hand, gradient attenuation is an emergent property born because it arises from the combination of two mechanisms: the reset mechanism and of the locality of the influence between pixels, determined by the sampling mechanism.

A last clarification: as to the thresholding mechanisms, even if it is not essential for gradient attenuation, it is an effective mechanism (as shown already in [22] page 2617 Fig 1) and one can use it to attenuate *small gradients*. The gradient attenuation property determined by reset and locality, on the contrary, operates on all (small and large) gradients.

## 5 ACKNOWLEDGMENTS

The authors acknowledge the support of the ICT Fund at EBTIC / Khalifa University, Abu Dhabi, UAE. The work was partially funded by the EU H2020 Research Programme, within the projects Threat-Arrest (Grant Agreement No. 786890) and Concordia (Grant Agreement No. 830927).

## REFERENCES

- [1] Banić, N., Lončarić, S.: Light random sprays retinex: exploiting the noisy illumination estimation. *Signal Processing Letters, IEEE* **20**(12), 1240–1243 (2013)
- [2] Frankle, J.A., McCann, J.J.: Method and apparatus for lightness imaging (May 17 1983), uS Patent 4,384,336
- [3] Funt, B., McCann, J., Ciurea, F.: Retinex in matlab. *Journal of Electronic Imaging* **13**(1), 48–57 (2004)
- [4] Gianini, G.: Statistical aspects of space sampling in retinex models. *Journal of Electronic Imaging* **20**(6), 1–6 (2016)
- [5] Gianini, G., Lecca, M., Rizzi, A.: A population-based approach to point-sampling spatial color algorithms. *JOSA A* **33**(12), 2396–2413 (2016)
- [6] Gianini, G., Manenti, A., Rizzi, A.: Qbrix: a quantile-based approach to retinex. *JOSA A* **31**(12), 2663–2673 (2014)
- [7] Gianini, G., Rizzi, A.: A Fuzzy Set Approach to Retinex Spray Sampling. *Multi-media Tools and Applications* **76**(23), 24723–24748 (2017)
- [8] Gianini, G., Rizzi, A., Damiani, E.: A retinex model based on absorbing markov chains. *Information Sciences* **327**, 149–174 (2016)
- [9] Kolås, Ø., Farup, I., Rizzi, A.: Spatio-temporal retinex-inspired envelope with stochastic sampling: A framework for spatial color algorithms. *Journal of Imaging Science and Technology* **55**(4), 40503–1 (2011)
- [10] Land, E.H., McCann, J.J.: *Lightness and Retinex Theory*
- [11] Lecca, M., Modena, C.M., Rizzi, A.: T-Rex: A Milano Retinex Implementation Based on Intensity Thresholding, pp. 68–79. Springer, Cham (2017)
- [12] Lecca, M., Rizzi, A.: Tuning the locality of filtering with a spatially weighted implementation of random spray retinex. *JOSA A* **32**(10), 1876–1887 (2015)
- [13] Lecca, M., Rizzi, A., Gianini, G.: Energy-driven path search for termite retinex. *J. Opt. Soc. Am. A (JOSA A)* (2016)
- [14] Lecca, M., Rizzi, A., Serapioni, R.P.: Grass: A gradient-based random sampling scheme for milano retinex. *IEEE Trans. Image Processing* **26**(6), 2767–2780 (2017)
- [15] Lecca, M., Rizzi, A., Serapioni, R.P.: Great: a gradient-based color-sampling scheme for retinex. *JOSA A* **34**(4), 513–522 (2017)
- [16] Marini, D., Rizzi, A.: A computational approach to color illusion. *Proc. of ICIAP97 Int. Conf. on Image Analysis and Processing* (1997)
- [17] Marini, D., Rizzi, A.: A computational approach to color adaptation effects. *Image and Vision Computing* **18**, 1005–1014 (2000)
- [18] Marini, D., Rizzi, A.: Colour constancy and optical illusions: a computer simulation with retinex theory. *ICIAP93 Int. Conf. on Image Analysis and Processing* (1993)
- [19] McCann, J., Rizzi, A.: *The Art and Science of HDR Imaging*. John Wiley (2011)
- [20] McCann, J.J.: Retinex at 50: color theory and spatial algorithms, a review. *Journal of Electronic Imaging* **26**(3) (2017)
- [21] Moore, A., Allman, J., Goodman, R.M.: A real-time neural system for color constancy. *IEEE Transactions on Neural networks* **2**(2), 237–247 (1991)
- [22] Provenzi, E., De Carli, L., Rizzi, A., Marini, D.: Mathematical definition and analysis of the retinex algorithm. *JOSA A* **22**(12), 2613–2621 (2005)
- [23] Provenzi, E., Fierro, M., Rizzi, A., De Carli, L., Gadia, D., Marini, D.: Random spray retinex: a new retinex implementation to investigate the local properties of the model. *Image Processing, IEEE Transactions on* **16**(1), 162–171 (2007)
- [24] Provenzi, E., Gatta, C., Fierro, M., Rizzi, A.: A spatially variant white-patch and gray-world method for color image enhancement driven by local contrast. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **30**(10), 1757–1770 (2008)
- [25] Richards, W., Parks, E.A.: Model for color conversion. *JOSA* **61**(7), 971–976 (1971)
- [26] Rizzi, A., McCann, J.: On the behavior of spatial models of color. *Proc. of Electronic Imaging 2007*. S. Jose, California (USA) (2007), (invited paper)
- [27] Rizzi, A.: Designator retinex, milano retinex and the locality issue. *Electronic Imaging* **2016**(6), 1–5 (2016)
- [28] Rizzi, A., Bonanomi, C.: Milano retinex family. *J. of Electronic Imaging* **26**(3), 031207–031207 (2017)
- [29] Rizzi, A., Marini, D., De Carli, L.: Lut and multilevel brownian retinex colour correction. *Int. J. Machine Graphics & Vision* **11**(2/3), 153–168 (2002)
- [30] Simone, G., Audino, G., Farup, I., Albrechtsen, F., Rizzi, A.: Termite retinex: a new implementation based on a colony of intelligent agents. *Journal of Electronic Imaging* **23**(1), 013006–013006 (2014)
- [31] Simone, G., Cordone, R., Serapioni, R.P., Lecca, M.: On edge-aware path-based color spatial sampling for retinex: from termite retinex to light energy-driven termite retinex. *Journal of Electronic Imaging* **26**(3), 031203–031203 (2017)
- [32] Sobol, R.: Improving the retinex algorithm for rendering wide dynamic range photographs. *Journal of Electronic Imaging* **13**(1), 65–74 (2004)
- [33] West, G., Brill, M.H.: Necessary and sufficient conditions for von kries chromatic adaptation to give color constancy. *J. of Math. Biology* **15**(2), 249–258 (1982)