QBRIX: a quantile-based approach to retinex

Gabriele Gianini, Andrea Manenti, and Alessandro Rizzi*

Universita degli Studi di Milano, Dipartimento di Informatica via Bramante 65, Crema (CR) 26013, Italy
*Corresponding author: alessandro.rizzi@unimi.it

In this paper, we introduce a novel probabilistic version of retinex. It is based on a probabilistic formalization of the random spray retinex sampling and contributes to the investigation of the spatial properties of the model. Various versions available of the retinex algorithm are characterized by different procedures for exploring the image content (so as to obtain, for each pixel, a reference white value), then used to rescale the pixel lightness. Here we propose an alternative procedure, which computes the reference white value from the percentile values of the pixel population. We formalize two versions of the algorithm: one with global and one with local behavior, characterized by different computational costs.   © 2014 Optical Society of America

1. INTRODUCTION

Retinex [1] is a model belonging to a wider family of algorithms, called spatial color algorithms [2]. When applied to an image to be enhanced, retinex filtering performs an automatic adjustment of the visual contrast of the input image. From the point of view of the image content, it operates as an edge enhancer and a suppressor of smooth gradients, in analogy to the human vision system.

Indeed, retinex was originally conceived as an abstract computational model of human color sensation; however, with time, it also has been used for very different purposes, as will be mentioned below. All the various retinex variants share the idea of recomputing the color of each pixel based on the spatial distribution of the other pixels' values in the image [2]. This trait makes the retinex behavior dependent on the image to be filtered: the algorithm adapts to the image content, thus equalizing the spatial distribution of contrast and yielding an image enhancement, which is robust with respect to the image variability.

Differences arise among the retinex algorithms depending on their purpose and application domains. Among typical purposes and applications are color constancy [3–5], separation of illumination from reflectance [6], shadow removal [7], HDR imaging [8,9], human vision modeling [1,10,11], photographic dynamic range rendering [12,13], color adjustment for pictures taken under unknown lighting conditions [14,15], and unsupervised color movie restoration [16].

Although the disparate problems to solve have given rise to different formulations and implementations [17], whose goals and success criteria are quite different for the different tasks, the success of retinex derives mainly from its specific local filtering properties. More specifically, when implemented as a *global* filter, retinex results in a simple Von Kries scaling [18]; however, as such, it can be useful only within a limited range of situations. It is the *local* behavior of retinex—i.e., the dependence of the pixel correction upon the spatial arrangement of the values around it—which makes it a useful model. From now on, we will refer to this feature by the term *locality*.

The various ways its locality is realized characterize the different implementations [17]. Among the many image sampling methods used, we recall predefined [19–21], constrained [22], or Brownian random paths (isotropic memoryless random walks) [11], fixed masks [23], random points [24], multilevel image decomposition [20,21,25]. Moreover, several variational formulations of the model also have been developed so far [6,10,12,26,27].

All the versions had to deal with the high computational cost necessary to realize some form of locality. The random spray retinex (RSR) algorithm [24] faced this issue by leveraging the presence of redundancy in the random path model [11]: to this purpose, it used a random point sampling schema in place of the Brownian walks. The model was endowed with better computational efficiency and, at the same time, with great versatility, since the exploration of the local behavior could be based on several point sampling profiles; as with all the sampling-based algorithms, however, in its application, one had to look, case by case, for the appropriate trade-off between impact of sampling noise and computational cost.

In this work, we introduce a probabilistic formulation of RSR, which is, by construction, free from sampling noise and provides even greater versatility in analyzing the retinex behavior, in that it is not bounded by the efficiency limitations of the sampling techniques and can adopt, with ease, virtually any function as a sampling profile.

The paper is organized as follows. First, we present some characteristics of the retinex model and describe how it samples the information from the image, with special reference to RSR (Section 2). Then we discuss the rationale of the QBRIX approach (Section 3) and the global QBRIX implementation (Section 4). Finally, we introduce the complete version, called local QBRIX, which directly compares with RSR (Section 5).

2. IMAGE SAMPLING IN RETINEX AND RSR

Retinex reflects an important basic property of our vision system: the *color sensation* at each point of the scene does not

derive solely from the light spectral characteristics at the point, but from the ratios—separately for each chromatic channel—of the channel values at the point to the values at other points in the scene.

In this paper, with the term *sensation*, we refer to the definition given by the committee on Colorimetry of the Optical Society of America in [28]. To distinguish between sensation and perception, it reported the following definitions. Sensation is the mode of mental functioning that is directly associated with stimulation of the organism. Perception is the mode of mental functioning that includes the combination of different sensations and the utilization of past experience in recognizing the objects and facts from which the present stimulation arises. Hence, those complex cognitive tasks, or even visual tasks, which require recognition of parts of the scene content—and which are at the base of many interesting visual illusions—do not fall within the scope of the term *sensation*.

A wide number of different retinex implementations have been realized so far. An overview is out of the scope of this paper; however, these implementations can be divided into two major groups and differ in the ways they achieve locality. The first group, among which we can mention RSR, [1,9,11,19–21] uses a sampling approach: the neighborhood of each pixel is explored either using paths or extracting random pixels; the second group [23,29–32] computes values over the image with convolution masks or weighting distances. An extensive review on retinex, including recent PDE and variational implementations, can be found in [33].

A. MI Retinex

In this paper, we start from the description of the Brownian random path retinex [11] called "MI retinex" [33] [from the city (Milano) where the algorithm has been developed and tested], which gave rise to a set of retinex versions including MLV retinex [25], RSR [24], and RACE [34].

In MI retinex, each random path is used to modify only the pixel at the end of the path (from now on called *target pixel*) as opposed to other versions, which modify all the pixels along a path, see, for example, [22]): the Brownian random path retinex computes each pixel value as the ratio between the original value of the pixel itself and the lightest pixel found by path random scanning across the image. This is done on each chromatic channel separately. We briefly recall the mathematical description of retinex given in [35].

Given an RGB digital image, consider a collection of N paths $\tilde{\gamma}_1, \ldots, \tilde{\gamma}_k, \ldots, \tilde{\gamma}_N$, consisting in ordered chains of pixels starting at pixel j_k and ending at pixel i (called *target pixel*). Let n_k be the number of pixels traveled by the kth path $\tilde{\gamma}_k$. We reserve the symbol ℓ to refer to a path index other than n_k, i.e., $\ell \in \{1, \ldots, (n_k - 1)\} \to \text{Image} \subset \mathbb{R}^2$, with $\tilde{\gamma}_k(1) = j_k$ and $\tilde{\gamma}_k(n_k) = i$.

We indicate two successive pixels in the path as $\tilde{\gamma}_k(\ell) = x_\ell$ and $\tilde{\gamma}_k(\ell + 1) = x_{\ell+1}$. In MI retinex, as in RSR, the computations are performed separately and independently in each chromatic channel $c \in \{R, G, B\}$ over the original image intensities $I_c(x)$ of the image pixels. For sake of simplicity, we will omit the chromatic channel specification. So we write $I(x_\ell)$, $I(x_{\ell+1})$ to indicate the input intensities of x_ℓ and $x_{\ell+1}$ (in each of the three chromatic channels, which we process separately). We write their ratio as $R_\ell = I(x_{\ell+1})/I(x_\ell)$. For

practical reasons, it is useful to rescale the possible values of intensities $I(x)$ by mapping their values in the real unit interval [0,1]. The output value of pixel i is the normalized *lightness $L(i)$*, which is computed as follows: when $I(i) = 0$ the output lightness is $L(i) = 0$; when $I(i) \in (0, 1]$ the output lightness $L(i)$ is given by the function:

$$L(i) = \frac{1}{N} \sum_{k=1}^{N} \left[\prod_{\ell=1}^{n_k-1} \rho_k(R_\ell) \right], \tag{1}$$

where $\rho_k : \mathbb{R}^+ \to \mathbb{R}^+$, $k = 1, \ldots, N$, is defined on the basis of the ratios R_ℓ as follows: $\rho_k(R_1) = 1$ whereas, for $\ell = 1, \ldots, n_k - 1$,

$$\rho_k(R_\ell) = \begin{cases} R_\ell & \text{if } 0 < R_\ell \leq 1 - \varepsilon \\ 1 & \text{if } 1 - \varepsilon < R_\ell < 1 + \varepsilon \\ R_\ell & \text{if } 1 + \varepsilon \leq R_\ell \leq \frac{1+\varepsilon}{r_k^{\ell-1}} \\ \frac{1}{r_k^{\ell-1}} & \text{if } R_\ell > \frac{1+\varepsilon}{r_k^{\ell-1}} \end{cases}, \tag{2}$$

where $\varepsilon > 0$ is a fixed *threshold* and $r_k^{\ell-1}$ represents the retinex's chain product of ratios for a given path $\tilde{\gamma}_k$ up to step $(\ell - 1)$ and is given by

$$r_k^{\ell-1} \equiv \prod_{m=0}^{\ell-1} \rho_k(R_m).$$

When the first or the third conditions are satisfied, ρ_k acts simply as the identity function, and Eq. (1) implements a simple average over chain product of ratios [1].

The second condition occurs when only a very small change of intensity is measured between two subsequent pixels. In this case $\rho_k(R_\ell) = 1$, so that the product of ratios remains unchanged with respect to the previous step. This implements the so-called threshold mechanism [1]. As will be later clarified, RSR does not implement the threshold mechanism [24] and, consequently, neither does QBRIX. For natural image enhancement, the absence of the threshold mechanism does not lead to remarkable differences and brings the advantage of a relative computational simplification.

Finally, when $\rho_k(R_1)\rho_k(R_2)\cdots\rho_k(R_{\ell-1})R_\ell > 1 + \varepsilon$, i.e., when the fourth condition is satisfied, ρ_k resets the chain of products to 1, so that the latest encountered pixel becomes the local reference white (the value with respect to which the pixel intensity will eventually be rescaled). This is the so-called reset mechanism.

Notice that the output of the algorithm, i.e., the lightness of pixel i, is a positive scalar, normalized to the interval $[0, I_{\text{TOP}}]$. With the above conventions about the input, we have $I_{\text{TOP}} = 1$.

All the options together realize the well-known ratio-threshold-reset mechanism of retinex. The white patch behavior (the correction toward a reference white value) of the algorithm is determined by this mechanism. A more detailed discussion can be found in [33].

B. Random Spray Retinex

A relatively recent implementation [24]—in order to reduce the sampling redundancy of the algorithm (Brownian paths can visit the same places many times) and in order to investigate the effects of different spatial samplings schemas—replaces paths with random sprays [hence, the name "random

spray retinex" (RSR)], i.e., 2D point distributions across the image, sampled according to a specific sampling profile.

RSR prescribes to correct the lightness of the target pixel i, from the original value $I(i)$ to a rescaled lightness value $L(i)$ defined by

$$L(i) = \frac{I(i)}{W(\Omega_i)}, \qquad (3)$$

where $W(\Omega_i)$ indicates the local reference white found by RSR computing a suitable average of the intensity maxima from a large number of n point samples, the *sprays* (see an example in Fig. 1).

More specifically, each spray picks n pixels at random from the target's neighborhood Ω_i, according to a radially symmetric sampling profile, then finds the pixel with maximum lightness and eventually uses its intensity as the value of its own contribution to the reference white, unless its intensity is lower than the original target intensity (in which case, the original intensity is used as a contribution). The latter rule is called the implicit reset mechanism and can be implemented by inserting by default the target pixel in the spray; hereafter, we will call bare spray the original n point spray and augmented spray the set, including the bare spray plus the target pixel. This sampling and selection procedure is repeated for a high number N of sprays; eventually, the reciprocal values of the sprays' contributions are averaged to obtain $W(\Omega_i)$.

We indicate by N the total number of sprays, by $\gamma_k \equiv \{x_1^k, x_2^k, \ldots x_l^k \ldots x_n^k\}$ the kth *bare* spray, by γ_k^* the kth *augmented* spray, defined by the union of the bare spray and the target pixel itself: $\gamma_k^* \equiv \{\{i\} \cup \gamma_k\}$, and, finally, by K_k the contribution of the kth augmented spray.

For each target pixel i, the detailed RSR steps are as follows:

1. Sample n points from the neighborhood Ω_i of the target pixel i according to a radially symmetric profile (see below and the implementation discussion for examples)

2. Add the target pixel to the set (to implement the implicit reset mechanisms) and obtain the $(n+1)$-points set γ_k^* defined above

3. Select the maximum H_k of the intensity found in γ_k^*:

$$H_k \equiv \max_{x \in \gamma_k^*} I(x)$$

4. Compute the contribution $K_k \equiv 1/H_k$

5. Repeat a number N of times the steps (1) through (4)

6. Compute the average $\bar{K}$ of the different contributions:

$$\bar{K} = \frac{1}{N} \sum_{k=1}^{N} K_k$$

7. Set $1/W(\Omega_i) = \bar{K}$ and rescale the intensity value of the target pixel from $I(i)$ to $L(i)$ by using Eq. (3), i.e., by $L(i) = I(i)\bar{K}$

Dependence on the local details of the target's closest neighborhood (i.e., locality) is provided to spray sampling by ascribing lower importance to pixels whose distance from i is greater: this is achieved by defining a radial weight function $w(d)$, where $d = d(i,j)$ indicates the Euclidean distance of the neighborhood pixel $j \in \Omega_i$ from the target pixel i, which determines the sampling profile (for instance, one can define $w(d) \propto 1/d(i,j)$). Figure 1 shows an example of an RSR spray and Fig. 2 its effect of its adoption on the RSR processing of an input image.

We observe, for further reference, that it is not possible to determine *a priori* an optimal configuration for RSR [functional form $w(d)$ of the sampling profile, number n of sample points in a spray, number N of sprays per pixel]: the optimal

Fig. 1.   Example of a "naturally localized" spray.

Fig. 2.   Example of RSR filtering. Upper image: the original image. Lower image: the output obtained using $N = 4$ natural sprays of $n = 150$ points each (the low value of N has been adopted only for demonstrative purposes, so as to highlight the intrinsic noisiness of the process).

configuration for a given image has to be found by an empirical trial and error process, i.e., an exploration of the configuration parameters' space.

3. QBRIX RATIONALE

The QBRIX algorithm is based on the following considerations about the original MI retinex and the RSR approximation.

Statistical sampling procedures are intrinsically noisy, and, as can be noticed in Fig. 2, the RSR one makes no exception: two neighboring target pixels with equal original lightness and equivalent neighborhoods can be assigned perceptibly different corrections because different samples have been used in the lightness computation. Mitigating this kind of noise requires the use of a large number of sprays per target pixel, i.e., increasing the computational cost. These drawbacks are common to all the sampling algorithms [2]. The probabilistic version we are going to introduce in the next section computes the quantity of interest [i.e., $W(\Omega_i)$] directly, based on the pixel population intensity distributions and, therefore, by construction, is not affected by sampling noise.

A. Key Idea: from Statistical Samples to Pixel Intensity Populations

The key idea underlying QBRIX is that the sampling procedure used by RSR can be replaced by an equivalent exact probabilistic computation.

Hereafter, we will elaborate on the idea that the RSR procedure can be reinterpreted as the statistical sampling estimate of a specific quantity and then exploit the fact that, in our case, there is no strict need of resorting to such an estimate, since the population is known: that specific quantity can be computed exactly based on the population. We will show that the value computed by the RSR procedure corresponds, in statistical terms, to an *estimator* for the mean of the sampling-minimum distribution of the inverse intensity, in short, from now on MSMDII (we will see that this, according to RSR, *is* the reference white value). QBRIX provides an alternative procedure, which computes directly the MSMDII, based on the high percentile values of the pixel intensity distribution.

In the following discussion, in order to arrive by steps to the probabilistic version of RSR, we will distinguish two main forms of spatial sampling: we will consider first the spatially homogeneous sampling of the neighborhood, which will originate a nonlocal, i.e., global, version of the algorithm (global QBRIX, Section 4); then we will consider the case of spatially nonhomogeneous sampling, which will correspond to the locality-aware version of the algorithm (local QBRIX, Section 5).

Let us start considering as a neighborhood of the target point i [with intensity $I(i)$] the whole image ($\Omega_i \equiv$ Image), and let us assume a spatially homogeneous sampling profile for RSR: in this case, the individual n point spray picks point uniformly at random within the image. The reference statistical population, in this case, is the set of pixel intensities of the whole input image. We can represent this population simply by the image intensity histogram. We will indicate by I the variable representing the intensity of the pixels in the image and by $f(I)$ the intensity distribution of this population.

The individual RSR n point spray draws a sample γ of n pixel intensities from this population (step 1), this is the *bare*

spray; let us indicate by H the maximum intensity value of the bare spray; then the intensity value of the target pixel is added to the set, so as to obtain the *augmented spray* γ^* (step 2); at this point, the maximum $H^* = \max\{I(i), H\}$ of such a new set is selected (step 3) and used to compute the reciprocal $K^* = 1/H^*$ (step 4), which is taken as the contribution of the spray to the reference white value of the target. We are interested in the sampling distribution of H^* and, consequently, of K^*. In fact, it is immediate to observe that repeating steps 1 through 4 and then averaging the obtained contributions (i.e., applying RSR) is equivalent to estimating the mean of the sampling distribution of $1/H^*$. However, if the sampling distribution of $1/H^*$ can be worked out, then its mean can be computed directly.

It turns out that such a distribution and its mean can indeed be determined, based on the image histogram $f(I)$, by using standard probability theory methods (as will be shown hereafter).

It also will be apparent shortly that the mean of the distribution of $1/H^*$ is not perceptibly different from the mean of the distribution of $1/H$ and that both means correspond to a high quantile of $f(I)$.

B. Tuning by Quantiles

The latter point brings about another key consideration: although it is possible to determine an *exact* mapping from the RSR parameters to the QBRIX parameters, such a mapping is not needed in practice. Given an image, it is experimentation that guides the choice of the optimal RSR configuration. Similarly, experimentation should guide the tuning of QBRIX.

The advantage in the latter case is that experimentation corresponds to a very simple process because it reduces to trying different high-level quantiles of the intensity distribution as reference white values for the target.

4. GLOBAL-QBRIX (SPATIALLY HOMOGENEOUS RSR SAMPLING)

It is rather straightforward to develop the probabilistic version of RSR in the case of *spatially homogeneous* sampling; that is, in the case where the neighborhood of the target point encompasses the whole image and the sampling profile of the spray, it is assumed to be constant, i.e., independent of the distance from the target. By construction, the corresponding algorithm operates on the input image a correction, which is insensitive to the local details of the target pixel neighborhood. In order to emphasize this aspect, we call this version of the algorithm global QBRIX.

Hereafter, we develop first an illustrative case, where the mapping from RSR to QBRIX can be computed exactly and where an analytical expression can be provided that links the number of points in the RSR spray to the quantile of the intensity distribution to be used as reference white level for the whole image.

This mapping will give the opportunity to discuss crucial details of the two algorithms (related to the *reset* mechanism) and to legitimate an approximation adopted in QBRIX.

A. Uniform Intensity Histogram

We will show now that, in the hypothesis that intensity levels are all equally represented in the input image, the use of an n

points RSR spray is equivalent to the adoption of the $(1 - 1/n)$ quantile as a reference white value for the image.

We recall that, given the intensity distribution $f(I)$ of the image and the corresponding cumulative $F(I)$, the inverse $F^{(-1)}$ of the cumulative, called *quantile function* and denoted by Q, maps a real number $p \in [0, 1]$ (representing a probability) to a value I of the variable; hence, $I = F^{(-1)}(p) = Q(p)$ and $Q{:}p \to Q(p)$. A value of a variable I corresponding to a given value of probability p is generically called p quantile and indicated by q_p.

We now initially assume, for illustrative purposes, that $f(I)$ is uniform (i.e., that all the lightness levels are equally represented in the neighborhood of the target pixel, equivalently. The corresponding histogram is flat) and therefore equal to 1 if I is defined in the unit interval [0,1] (adopted following the conventions introduced previously); this yields the cumulative:

$$F(I) = \begin{cases} 0 & \text{if } I \le 0, \\ I & \text{if } 0 < I \le 1, \\ 1 & \text{if } I > 1. \end{cases}$$

We focus on the behavior within the interval $I \in [0, 1]$, where $F(I) = I$.

1. Bare Spray (No Implicit Reset)

Let us now reformulate RSR steps within this probabilistic framework. Initially, for sake of simplicity, we skip step 2 (corresponding to the addition of the target intensity to the sample set), i.e., work out the correspondence by assuming the use of bare sprays (no implicit reset); however, we are going to return shortly to this point.

Here are the RSR steps and the corresponding probabilistic interpretation:

Steps 1 and 3: sample from $f(I)$ a set $\gamma = \{x_1, x_2, \dots x_l \dots x_n\}$ of n points and pick the value H corresponding to the maximum $H = \max_{x_j \in \gamma} I(x_j)$.

The cumulative $F_{\max}$ of the probability density of the maximum H of n points sampled from $f(I)$ is given by the cumulative $F(.)$ raised to the power of n; hence,

$$F_{\max}(H) = F^n(H) = H^n,$$

(for later reference, we indicate the density of such a maximum by $f_{\max} = dF_{\max}/dH$).

Step 4: Compute $K = 1/H$:

when we perform a variable change $H \to K = 1/H$ the density defined in the interval $H \in (0, 1]$ changes into a density $g(K)$ defined for $K \in [1, +\infty)$, such that $F(H) = \Pr(x < H) = \Pr(1/x > K) = S(K)$; hence,

$$S(K = 1/H) = F_{\max}(H) = H^n = 1/K^n,$$

where $\Pr(\cdot)$ is the probability of the event.

Steps 5 and 6: Perform N repetitions of the sampling and take the average.

We are interested in the expected value of the density $g(K)$. The average of K can be obtained in this case based on the density $g(K) = -dS/dK$ or from the complementary cumulative method:

$$\bar{K} = \int_1^{+\infty} K g(K) dK = \int_1^{+\infty} K \left(-\frac{dS}{dK} \right) dK$$

$$= [K S(K)]_1^{+\infty} + \int_1^{+\infty} S(K) dK = 1 + \frac{1}{n-1}.$$

It turns out that the reference white value is simply

$$W(i) = \frac{1}{\bar{K}} = 1 - \frac{1}{n}.$$

Step 7 (for RSR and QBRIX): Use this average to obtain the lightness value of the target pixel as

$$L(i) = I(i)/W(i) = I(i) \cdot \bar{K}. \tag{4}$$

Notice that the reference white value corresponds to the $(1 - 1/n)$ quantile of the original variable I. For instance, for a value of $n = 100$ points in the RSR spray, the expected value of the reference white value falls on the 99th percentile of the variable I. This straightforward relation holds, thanks to the fact that we are assuming a uniform intensity distribution (i.e., a flat intensity histogram). In the case of a nonuniform density, the quantile position will be different and depend mostly on the characteristics of the upper end of the lightness distribution. Notice, furthermore, that for the pixels with original intensity greater than the reference quantile, the application of the relation in Eq. ([4](#)) would yield lightness values greater than 1. This is due to the lack of the implicit reset in the bare spray formulation. Now we discuss this point.

2. Augmented Spray (Implicit Reset)

Now we consider how the expected value changes when using step 2, i.e., when also inserting the intensity $I_i \equiv I(i)$ of the target in the spray so as to obtain an augmented spray. In this way, no spray maximum can fall below I_i. Including i in the spray corresponds to censoring inferiorly the distribution of H at the threshold point I_i and transferring the excess mass from below the threshold onto the threshold point; equivalently, it corresponds to censoring superiorly K at the threshold point $1/I_i$ and transferring the excess mass onto the threshold point. Recomputing the mean, we find out that the use of the *augmented* spray, in place of the *bare* spray, introduces the following correction to the final result:

$$\bar{K}^* = \bar{K} \times \left(1 - \frac{I_i^{n-1}}{n} \right). \tag{5}$$

Let us consider two distinct examples with reference to the case of the $n = 100$ points spray and, consequently, the case $W(i) = 1/\bar{K}$ corresponding to the 99th percentile of the variable I. We will consider one target pixel with intensity below this threshold (at $I_i = 0.95$) and one with intensity above this threshold (with $I_i = 0.995$).

If the intensity of the target pixel is $I_i = 0.95$ (i.e., in our uniform distribution of intensity case, is at the 95th percentile), the relative change from $\bar{K}$ to $\bar{K}^*$—thus, in the computed output lightness—is of the order of 10^{-4}. Thus, in place of $\bar{K}^*$, we can safely use $\bar{K}$ (and the corresponding quantile $1/\bar{K}$ as a reference white value level).

If the intensity of the target pixel is $I_i = 0.995$, the use of the 99th percentile as a reference white value would yield illegal

output values of the lightness greater than 1. Using the correct expression in Eq. (5) (notice the negative sign of the correction) issues an $L(i) = I(i)\bar{K}^*$ just below 1. The important observation here is that, however, the difference is so tiny that one can safely use $L(i) = 1$ at its place. For an intensity $I_i = 0.995$: using $\bar{K}^*$ yields an output lightness $L(i) = I(i)\bar{K}^*$ such that $1 - L(i) \simeq 10^{-3}$. This prompts for setting to 1 the lightness of all the points exceeding the reference quantile.

It can be shown that the above two approximations are rather accurate, in general, for pixels with intensities respectively below and above the threshold.

The bottom line is the following: when willing to reproduce the effect of a large collection of n point RSR sprays (under uniform distribution of intensity hypothesis), we can use the approximation $W(i) = 1/\bar{K}^* \simeq 1/\bar{K} = 1 - 1/n$, i.e., adopt the $(1 - 1/n)$ quantile of the original variable as a reference white level; if the original intensity of the target exceeds such a reference value, we set its lightness to $I_{\text{TOP}} = 1$.

This is the essence of the QBRIX approach. This result concludes the discussion of our illustrative example based on (homogeneous spatial sampling and) uniform intensity histograms.

B. Generic Intensity Histograms
Assuming the uniform intensity histograms we have shown, how, in place of using the RSR sprays, can one compute the mean of the sampling minimum distribution of the inverse intensity (MSMDII) in terms of quantiles of the original intensity distribution? We also have shown that the quantile value representing the equivalent of a specific RSR spray prescription, relates to the number n of points in the corresponding RSR spray.

This qualitative relation holds also when computing the MSMDII under more realistic (nonuniform) intensity distributions. This can be shown by retracing the derivation of the previous subsection from a generic (nonuniform) distribution of intensities $f(I)$ all the way to the corresponding average $\bar{K}$.

For better compactness, notice that looking for the (mean of the) sampling distribution of the reciprocal of the maximum is the same as looking for the (mean of the) sampling distribution of the minimum of the reciprocal. From the distribution (histogram) $f(I)$ defined in $[0,1]$ and its cumulative $F(I)$ (or directly from the image pixel intensity), one can obtain the complementary cumulative $S(\mathcal{R})$ of the distribution of its reciprocal $\mathcal{R} = 1/I$, defined in $[1, +\infty)$. The complementary cumulative of such a minimum is $S_{\min}(\mathcal{R}) = S^n(\mathcal{R})$ for a sample of n points, rewritten in terms of the minimum variable K, $S^n(K)$ again defined in $[1, +\infty)$. Exploiting, as before, the complementary cumulative method to compute the mean, we have (first equality)

$$\bar{K} = 1 + \int_1^{+\infty} S^n(K)\mathrm{d}K \simeq 1 + \sum_{K=1}^{K_{\text{tr}}} S^n(K).$$

In the second equality, K runs over the histogram bins; considering that a high power n of values lower than 1 makes them noninfluential, the sum can be truncated after a limited number of steps K_{tr} without affecting the result.

Although the above expression provides an essentially *exact* mapping from RSR to QBRIX based on the RSR parameter n and the intensity histogram, it will not be used

in practice when processing an image: experimentation with different high-order quantiles will guide the choice of the optimal QBRIX configuration, as will be exemplified shortly.

This consideration also will hold for nonhomogeneous sampling profiles (see local QBRIX, next section) in which case the value of the quantile to be adopted also will depend on the shape of the sampling profile. For this reason, we will not retrace the updated formal derivation in that case.

C. Implementation of Global QBRIX
As all the other retinex implementations, following the original retinex approach, QBRIX operates independently in the three chromatic channels.

The algorithm to find a unique reference white value for the whole image (corresponding to a spatially homogeneous sampling in RSR) is the following.

For each chromatic channel $c \in \{R, G, B\}$ (the indication of the chromatic channel will be omitted for notational simplicity), take the following steps:

1. Get, from the whole pixel set of the image, the overall distribution $f(I)$, corresponding to the image histogram, and from this obtain the cumulative $F(I)$

2. Choose a high value of $p \in (0, 1]$ and find the corresponding p quantile q_p based on the inverse of the cumulative distribution $F(I)$

3. For every pixel i,

– if $I(i) < q_p$ rescale the value of the pixel from $I(i)$ to $I(i)/q_p$,
– otherwise, reset $I(i)$ to I_{TOP}.

In other words, the output lightness value of each pixel i according to global QBRIX can be found by this formula:

$$L(i) = \begin{cases} \frac{I(i)}{q_p}, & I(i) \le q_p \\ I_{\text{TOP}}, & I(i) > q_p \end{cases}.$$

We recall that image pixel values are normalized, so I_{TOP} is set to 1.

Notice, cursorily, that—since RSR estimates the reference white value by means of a sampling and averaging procedure—the number of sprays N is part of the RSR algorithm definition: a higher N value implies a more accurate estimate of the average value $\bar{K}$. On the contrary, QBRIX does not use a sampling procedure; therefore, N is not part of its definition.

D. Results, Intensity Histograms, and Locality Issues
As in RSR, where the parameter n affects the quality of the corrected image, in QBRIX different (high-order) quantiles of the lightness distribution result in filtering differences. Figure 3 shows a comparison among results from different quantile values used as reference white values. One can notice how the algorithm works differently according to the image content: darker pictures are more sensitive to quantile variations. From a series of tests on a wider set of images, we have found that a quantile between 95 and 99 leads to the most satisfactory results.

This implementation of the global QBRIX presents all the features of RSR such as, for instance, the property of color dominant removal, visible in Fig. 4, but without the spray sampling noise, as visible in the comparison of Fig. 5.

Fig. 3.   Comparison among results from different quantile values. The first image of each row is the original; then, from left to right, the QBRIX results using the 99th, 97th, 94th, and 88th quantile as the reference white value.

The computational complexity of the algorithm is $O(M)$, where M is the number of pixels of the image to be filtered.

Notice that the images used for the test have a nonuniform intensity histogram, as most real images. Nonetheless, in those images, as in all the other images tested so far, QBRIX displays a satisfactory performance. This can be due to the fact that different distributions can lead to the same quantile value. Consider also that it is rare to find natural images with a spiky histogram, which would let emerge the quantile computation problems known in data manipulation (empty regions and spikes, which can create jumps and ambiguities in the quantile function).


Fig. 4.   QBRIX property of color dominant removal. Left: Original image. Right: Image produced by QBRIX with the 93rd quantile as the reference white level).

Fig. 5.   Comparison between the effects of RSR and QBRIX on the same image. The original (top) followed by the images produced with RSR (middle) and with QBRIX (bottom). The last image used the 99th quantile as the reference white level. Notice the overall similarity between the two filters but also the different noise level inserted in the case of RSR due to the low value of N (here $N = 8$).

5. LOCAL QBRIX (SPATIALLY NONHOMOGENOUS RSR SAMPLING)

An important issue remains to be faced: the locality of filtering. Using a single scalar value to adjust all the pixels of the image, results in a global filtering. However, it is well known that the properties of the areas surrounding a point in a scene can change the appearance of the point. The locality issue can be disregarded only when the global chosen quantile remains the same in every random sampled subportion of the image (locality). In other words, if the local cumulative's upper portions are the same as the global one. However, this case is not common for natural images.

Thus we developed a local version of the algorithm, named local QBRIX.

The leading principle has been to use RSR again as a guide and to specify local QBRIX along the lines of RSR probabilistic formalization; in RSR, the spray maximum should be taken from the samples obtained following the sampling profile, which assigns to pixels a probability of being selected that decreases with the distance from the target. The technique used to achieve in QBRIX an effect analogous to a nonhomogenous sampling has been based on a reweighting of the space variables, which brings the problem back to a form equivalent to homogenous sampling. We used a *weighting profile* with the same functional form of the *sampling profile*.

The desired result is achieved by computing the QBRIX quantile on the basis not of the actual histogram but on a histogram built by weighted contributions: all image pixels are still represented in the histogram population; however, each contributes with a weight depending on its distance from the target pixel.

A. Implementation of Local QBRIX

Local QBRIX uses a quantile approach, as QBRIX, for identifying the reference white level; however, instead of adopting the same intensity histogram $f(\cdot)$ for all the pixels of the image, we use a different histogram $f_i(\cdot)$ for each target pixel i. That histogram is a reweighted version of the original image histogram $f(\cdot)$, where each pixel $j \in$ Image enters the histogram bin, corresponding to its intensity $I(j)$, weighted by a weight $w(i,j)$, depending on the Euclidean distance $d(i,j)$; thus, indicating the histogram channel by the variable $\mathbb{I}$:

$$f_i(\mathbb{I}) = \frac{1}{\mathcal{W}_i} \sum_{j \in \text{Image}, I(j) = \mathbb{I}} I(j) w(i,j) \quad \forall \; \mathbb{I} \in [0, \text{MCQ}],$$

where $\mathcal{W}_i = \sum_j w(i,j)$ is the normalization term given by the sum of the weights, and where MCQ is the maximum channel quantization (in the histograms of the present paper MCQ $= 255$).

Among the disparate sampling profiles available for RSR, we decided to use, in this paper, the following:

$$w(i,j) = (d(i,j)/D)^{-a}, \tag{6}$$

where a is a positive scalar to tune locality and D is the image diagonal length. In this way, the locality is independent from the absolute size of the image. The parameter a tunes the locality: the higher a, the more local is the behavior of the algorithm.

Once the local histogram has been constructed at each point, the local QBRIX quantile identification process is the same as described in the previous section: one choses a quantile value and uses the *same quantile* for all the target pixels in the image. The fact that the weighting profile determines a different histogram for each target pixel grants that the reference white value corresponding to the quantile is local and potentially different for each pixel.

The computational complexity of local QBRIX is greater than that of global QBRIX, since the operations of histogram construction, cumulative distribution calculation, and p quantile determination are repeated for each pixel of the image (while in global QBRIX, they were executed only once). Local

Fig. 6. Mach bands pattern (input top left, output top right) as rendered by local QBRIX with the sampling profile of Eq. (6) and $a = 2$. The bottom graphs represent the corresponding scan lines.

Fig. 7. Comparison among results with different a values [see Eq. (6) and same percentile (99th)]. The first image is the original; then $a = 1, 2, 3$.

Fig. 8. Comparison among results with different percentiles. The images are obtained from the original image (leftmost image in Fig. 7) by applying local QBRIX with (from left to right) the 99th, 98th, 97th, and 96th percentile as a reference white value and the sampling profile of Eq. (6) with $a = 2$.

QBRIX has a computational complexity of $O(M^2)$, since when processing a target pixel, it has to visit anew all the M points of the image.

B. Results, Algorithm Locality, and Choice of Quantile

To demonstrate the local behavior of local QBRIX, we have applied it to a classic low-level visual illusion: the results agree qualitatively with the effect of these illusion on our vision system, as is visible in the scan lines of Fig. 6. This result can be obtained only by a local behavior.

From tests made on different spatial color algorithms [2], it resulted that a quadratic distance in Eq. (6) gives the most

natural results [36]. We have obtained similar results from the same test with local QBRIX; an example is presented in Fig. 7.

Table 1. Correspondence between Parameters in RSR and Local QBRIX

RSR	Local QBRIX
N	(no correspondence)
n	reference quantile
sampling profile	pixel weighting

Fig. 9.   Comparison between RSR (center, $N = 20$) and local QBRIX (right) with the same figure of sampling. The original is the leftmost image.

Figure 8 presents the same image of Fig. 7 filtered by local QBRIX using different percentile values from 99 to 96, with $a = 2$.

According to the basic quantile mechanism in QBRIX, the lower the quantile, the lighter the resulting image. From the visual comparison of Figs. 7 and 8, it can be noticed that changes in the percentile value do not affect the locality, while the distance (parameter a) does.

C. RSR Versus Local QBRIX Locality

In RSR, three basic parameters affect the final result:

• The number of sprays N controls the noise: increasing N lowers the chromatic noise. This has no correspondent parameter in local QBRIX.

• The number of points per spray n: increasing n produces a finer sampling. As explained in Section 3.A, in local QBRIX this corresponds to the quantile choice. For instance, assuming lightness levels equally represented and homogeneous spatial sampling, the use of an n points spray is equivalent to the adoption of the $(1 - 1/n)$ quantile as a reference white value.

• The locality of filtering is controlled by the sampling profile in RSR and local QBRIX. (For instance, in RSR the uniform sampling of the distance and in the radial direction produces the so-called *naturally localized spray* [24], inversely proportional to r^2.) In local QBRIX, the sampling profile is realized through the pixel weighting.

Table 1 summarizes the correspondence between the RSR and the local QBRIX distinctive features.

In order to compare local QBRIX with RSR, we have tested the two approaches with the same spatial sampling.

For this comparison, we have used a sampling profile more concentrated near the pixel to enhance the locality. For RSR, it has been obtained by generating a random point in polar coordinates, with the following formula:

$$\Delta x = \rho^2 \, \cos(\theta)$$
$$\Delta y = \rho^2 \, \sin(\theta), \qquad (7)$$

Fig. 10.   Details from Fig. 9. Left: RSR. Right: local QBRIX.

where $\rho \in \mathrm{RAND}[0, R]$ and $\theta \in \mathrm{RAND}[0, 2\pi]$.

For the corresponding filtering of the image with local QBRIX, we used the pixel weighting according to Eq. (6) with $a = 2$; the percentile used is the 99th.

As noted in Fig. 9, the two algorithms produce very similar images: the most noticeable difference is the sampling noise of RSR, which has better visibility in the enlarged details of Fig. 10.

6. CONCLUSIONS

In this paper, we have presented a novel version of retinex based on a probabilistic reformulation of random spray retinex spatial sampling. In RSR, a random sample of pixels is drawn from the target pixel neighborhood, and the maximum lightness is selected; then the reciprocals of those sample maxima are averaged over several samples, so as to obtain

the reference white value for the target. Even if more efficient than Brownian random paths, the RSR spray implementation of retinex can suffer from noise. A higher number of sprays is beneficial to RSR, as it allows for more accurate determination of the reference white value; however, practically this requires a great computational effort. The proposed algorithm is, by construction, noise free, with a computational effort of $O(M)$ for global QBRIX and $O(M^2)$ for local QBRIX, where M is the number of pixels in the image.

QBRIX has come from the idea of implementing RSR in a sampling noise free form by considering the behavior of RSR from a probabilistic point of view. We argued that the value computed by the RSR procedure corresponds to a high quantile of the distribution of the intensities of the image. When using the same quantile value as a reference white value for the whole image, this approach is computationally efficient and not affected by chromatic noise. Nevertheless, the price to pay to obtain such an efficient algorithm is the lack of local filtering. To overcome this limit, we have presented a local version, named local QBRIX, which creates a suitable probability distribution at every pixel and takes a high quantile as a reference white value for that pixel.

These two probabilistic formalizations aim to contribute to the understanding of the spatial properties in human vision system modeling, which can be used for deeper investigation of the retinex family of algorithms.

ACKNOWLEDGMENTS

REFERENCES

1. E. Land and J. McCann, "Lightness and retinex theory," J. Opt. Soc. Am. **61**, 1–11 (1971).
2. A. Rizzi and J. J. McCann, "On the behavior of spatial models of color," in *Proceedings of Electronic Imaging*, San Jose, California (USA), 2007.
3. D. Marini, A. Rizzi, and M. Rossi, "Color constancy measurement for synthetic image generation," J. Electron. Imaging **8**, 394–403 (1999).
4. J. J. McCann, S. McKee, and T. Taylor, "Quantitative studies in retinex theory: a comparison between theoretical predictions and observer responses to color mondrian experiments," Vis. Res. **16**, 445–458 (1976).
5. J. J. McCann, "Local/global mechanisms for color constancy," Die Farbe **34**, 275–283 (1987).
6. R. Kimmel, M. Elad, D. Shaked, R. Keshet, and I. Sobel, "A variational framework for retinex," Int. J. Comput. Vis. **52**, 7–23 (2003).
7. G. Finlayson, P. M. Hubel, and S. Hordley, "Color by correlation," in *Proceedings of Fifth Color Imaging Conference*, 1997.
8. C. Gatta, A. Rizzi, and D. Marini, "Perceptually inspired hdr images tone mapping with color correction," Intern. J. Imag. Sys. Tech. **17**, 285–294 (2007).
9. L. Meylan and S. Susstrunk, "High dynamic range image rendering using a retinex-based adaptive filter," IEEE Trans. Image Process. **15**, 2820–2830 (2006).
10. M. Bertalmio and J. Cowan, "Implementing the retinex algorithm with Wilson–Cowan equations," J. Physiol. Paris **103**, 101–119 (2009).
11. D. Marini and A. Rizzi, "A computational approach to color adaptation effects," Image Vis. Comput. **18**, 1005–1014 (2000).
12. M. Bertalmio, V. Caselles, and E. Provenzi, "Issues about retinex theory and contrast enhancement," Int. J. Comput. Vis. **83**, 101–119 (2009).
13. R. Sobol, "Improving the retinex algorithm for rendering wide dynamic range photographs," in *Proceedings of Electronic Imaging*, San Jose, California (USA), 2002.
14. G. Ciocca, D. Marini, A. Rizzi, R. Schettini, and S. Zuffi, "Retinex preprocessing of uncalibrated images for color based image retrieval," J. Electron. Imag. **12**, 161–172 (2003).
15. A. Rizzi, L. Rovati, D. Marini, and F. Docchio, "Unsupervised corrections of unknown chromatic dominants using a Brownian paths based retinex algorithm," J. Electron. Imaging **12**, 431–441 (2003).
16. A. Rizzi, A. J. Berolo, C. Bonanomi, and D. Gadia, "Unsupervised digital movie restoration with spatial models of color," Multimedia Tools Appls., doi:10.1007/s11042-014-2064-5 (posted online July 1, 2014).
17. J. J. McCann, ed., "Special session on retinex at 40," J. Electron. Imaging **13**, 6–145 (2004).
18. J. von Kries, "Sources of color science," in *Chromatic Adaptation*, D. L. MacAdam, ed. (MIT, 1970), pp. 109–119.
19. T. J. Cooper and F. A. Baqai, "Analysis and extensions of the Frankle–McCann retinex algorithm," J. Electron. Imaging **13**, 85–92 (2004).
20. J. Frankle and J. J. McCann, "Method and apparatus of lightness imaging," U.S. patent 4384336 (May 17, 1983).
21. B. Funt, F. Ciurea, and J. J. McCann, "Retinex in MATLAB," J. Electron. Imaging **13**, 48–57 (2004).
22. R. Montagna and G. Finlayson, "Constrained pseudo-Brownian motion and its application to image enhancement," J. Opt. Soc. Am. A **28**, 1677–1688 (2011).
23. D. J. Jobson, Z. Rahman, and G. A. Woodel, "Properties and performance of a center/surround retinex," IEEE Trans. Image Process. **6**, 451–462 (1997).
24. E. Provenzi, M. Fierro, A. Rizzi, L. De Carli, D. Gadia, and D. Marini, "Random spray retinex: a new retinex implementation to investigate the local properties of the model," IEEE Trans. Image Process. **16**, 162–171 (2007).
25. A. Rizzi, D. Marini, and L. De Carli, "LUT and multilevel Brownian retinex colour correction," Mach. Graph. Vis. **11**, 153–168 (2002).
26. M. Elad, R. Kimmel, D. Shaked, and R. Keshet, "Reduced complexity retinex algorithm via the variational approach," J. Visual Commun. Image Represent. **14**, 369–388 (2003).
27. M. Elad, *Retinex by Two Bilateral Filters*, Vol. **3459** of Lecture Notes in Computer Science (Springer, 2005), , pp. 217–229.
28. "The science of color," OSA Committee on Colorimetry Tech. Rep. (Optical Society of America , 1953).
29. K. Barnard and B. Funt, "Investigations into multi-scale retinex," in *Proceedings of Colour Imaging in Multimedia*, Derby, UK, 1998.
30. J. D. Cowan and P. C. Bressloff, "Visual cortex and the retinex algorithm," in *Proceedings of SPIE Human Vision and Electronic Imaging VII*, San Jose, California (USA), 2002, Vol. **4662**.
31. S. O. Huck and C. L. Fales, "Visual communication with retinex coding," Appl. Opt. **39**, 1711–1730 (2000).
32. E. Land, "Recent advances in retinex theory and some implications for cortical computations: color vision and the natural image," Proc. Natl. Acad. Sci. USA **80**, 5163–5169 (1983).
33. J. J. McCann and A. Rizzi, *The Art and Science of HDR Imaging* (Wiley, 2011).
34. E. Provenzi, C. Gatta, M. Fierro, and A. Rizzi, "A spatially variant white patch and gray world method for color image enhancement driven by local contrast," IEEE Trans. Pattern Anal. Mach. Intell. **30**, 1757–1770 (2008).
35. E. Provenzi, L. De Carli, A. Rizzi, and D. Marini, "Mathematical definition and analysis of the retinex algorithm," J. Opt. Soc. Am. A **22**, 2613–2621 (2005).
36. A. Rizzi, C. Gatta, and D. Marini, "A new algorithm for unsupervised global and local color correction," Pattern Recogn. Lett. **24**, 1663–1677 (2003).