Michela Lecca, **STAR: A Segmentation-based Approximation of point-based sampling Milano Retinex for Color Image Enhancement,** IEEE TRANSACTIONS ON IMAGE PROCESSING**,** Volume: 27, Issue: 12, Dec. 2018, DOI: 10.1109/TIP.2018.2858541

The final published version is available online at: https://ieeexplore.ieee.org/document/8417897

STAR: A Segmentation-based Approximation of point-based sampling Milano Retinex for Color Image Enhancement

Michela Lecca[^]

***Abstract*—Milano Retinex is a family of spatial color algorithms inspired by Retinex and mainly devoted to the image enhancement. In the so-called point-based sampling Milano Retinex algorithms, this task is accomplished by processing the color of each image pixel based on a set of colors sampled in its surround. This paper presents STAR, a *S*egmen*t*ation based *A*pproximation of the point-based sampling Milano *R*etinex approaches: it replaces the pixel-wise image sampling by a novel, computationally efficient procedure that detects once for all the color and spatial information relevant to image enhancement from clusters of pixels output by a segmentation. The experiments reported here show that STAR performs similarly to previous point-based sampling Milano Retinex approaches, and that STAR enhancement improves the accuracy of the the well known algorithm SIFT on the description and matching of pictures captured under difficult light conditions.**

***Index Terms*—Retinex, Milano Retinex, Spatial Color Algorithms, Image Enhancement**
EDICS: ELI-COL, TEC-RST, SMR-HPM

I. Introduction

MILANO Retinex is a family of spatial color algorithms mainly devoted to color image enhancement [1]. As suggested by the name, these algorithms are inspired by the Retinex theory [2], introduced by Edwin Land and John McCann to estimate the so-called human *color sensation*, i.e. the color as seen by humans. Before Retinex, the human color vision system was supposed to work similarly to a camera, where the color signal from any observed point correlates with the luminance of that point and is independent of the other colors present in the scene [3], [4]. The experiments carried out by Land and McCann showed that, on the contrary, the human color vision is a *local spatial process*. Precisely, the human color sensation depends not only on the photometric cues of the observed point, but also on the spatial arrangement of the surrounding colors, with the colors closer to the observed point influencing more the color sensation than those located farther.

Based on these principles, Land and McCann proposed an algorithm to estimate the human color sensation, starting from an RGB image. The image is first *pre-calibrated*, i.e. its digital values are converted into the actual luminances of the observed scene [5], [6]. The color channels of the pre-calibrated image are processed separately, and each of them is mapped on to a new image, called *channel lightness*. In order to collect information about the spatial color distribution, the image support, i.e. the set of pixels composing the image, is scanned by a number of random paths. These define the locality of the algorithm, i.e. the range of action of the color influence. Over each path, Retinex computes the product of the ratios of the intensities of adjacent pixels, preventing, of course, division by zero. Every ratio close to one is cast to one, and every cumulative product exceeding one is cast to one. Then, Retinex computes the lightness at each pixel as a moving average of the cumulative products over the paths scanning the image. This mechanism is called ratio-product-threshold-reset-average mechanism from its computational phases. Finally, the *color lightness*, i.e. the RGB image composed of the channel lightnesses, is remapped according to a reference scale of appearance. The resulting image is an estimate of the human color sensation.

When no pre- and post-calibrations are applied, the Retinex algorithm works as an *image enhancer*, i.e. it takes as input a RGB image and outputs a novel RGB image, that usually has a higher brightness and more visible details than the input.

Milano Retinex is a family of algorithms performing spatial color processing procedures in line with the principles of the Retinex lightness computation. These algorithms have been mainly used for the image enhancement purpose. In this framework, they take as input a RGB image, skip the pre- and post- calibration steps (which are fundamental for human color sensation), and process the color channels separately by implementing novel, alternative definitions of the locality and new equations for the lightness. In Milano Retinex, each pixel is regarded as a *target*, and the channel lightness is computed pixelwise from a set of visual features extracted from pixels around the target. The first Milano Retinex implementations, e.g. [7], scan the neighborhood of each target by a set of random paths ending at the target. The lightness is the average of the cumulative products of the intensity ratios along each path, within the threshold- and reset rules of the original Retinex. According to [8], subsequents Milano Retinex implementations replace the ratio-threshold-reset rule with an *intensity re-scaling procedure*. This latter computes the target channel lightness as the ratio of the target channel intensity and a *local reference white*, i.e. an intensity level depending on the intensities of the pixels around the target, and, in some cases, also on image gradient and spatial cues.

[^]M. Lecca is with Fondazione Bruno Kessler, Center for Information and Communication Technology, via Sommarive 18, Trento 38123, Italy, e-mail: lecca@fbk.eu. This paper has supplementary downloadable material available at http://ieeexplore.ieee.org, provided by the author. The material includes the full-resolution version of some images of this paper. This resolution is suitable to better appreciate image details. Contact lecca@fbk.eu for further questions about this work.

Manuscript received ...; revised ...

Many Milano Retinex approaches model locality through a spatial sampling procedure that analyzes the channel intensities over the image support [9] or over a part of it, such as paths [10], [11], [12], random sprays [13], [14], [15], edges [16], [17]. The local reference white at any target is computed by averaging the maximum of the intensity (e.g. [10], [11], [13], [14], [16]) or the maximals of the target intensity (e.g. [17] and [9]) over the sampled sets[1]. In some algorithms, these averages are weighted by functions of the distance between the sampled pixels and the target, e.g. [15], [16], [17], according to the empirical evidence that colors close to the target are more relevant than those far away. Other Milano Retinex algorithms, like [18], [19], [20], avoid the sampling procedure, and estimate the local reference white by a statistical analysis. In particular, the work in [18] re-writes the random paths sampling scheme in terms of absorbing Markov chains, while the papers in [19] and [20] propose a statistical model of the sampling scheme of [13].

The level of the image enhancement produced by the Milano Retinex methods depends on the locality model and on the lightness equation. In this respect, the Milano Retinex family provides a large variety of image enhancement techniques.

This paper presents STAR, a novel spatial color algorithm for image enhancement, inspired by the Milano Retinex approaches. STAR is a *S*egmen*t*ation based *A*pproximation of Milano *R*etinex sampling based methods. STAR aims at cutting down the computational burden of the image sampling and at reducing the number of operations needed to compute the lightness. STAR accomplishes this task by employing *coarse* color and distance information, computed from clusters of pixels detected by a segmentation and independent of the target. Moreover, STAR models locality by considering the mutual distances between the segments, instead of the pixel-wise distances between the target and the sampled pixels, as usually done in many other Milano Retinex approaches.

STAR takes as input a RGB image and processes its color channel independently. The lightness of each channel is computed by two steps. In the first step, STAR performs a *global* channel analysis. STAR segments the channel in many regions, let's say $P_1, \ldots, P_M$, with $M > 0$. For each region P_j, STAR selects a set of *representative elements*, which are the pixels maximizing the distance transform over P_j, and computes the maximum intensity value over P_j. Then, it computes the mutual distances between the segmented regions as the minimum distances between their representative elements. In the second step, STAR implements a *pixel-wise* analysis. For each target x, STAR detects the segment P_i including x, then computes the lightness at x as the average of the maximals of the target intensity over the M intensity values selected from the channel segments. The contribution of every maximal to the lightness is weighted by a function inversely proportional to the distance between P_i and the region containing the maximal, defined in the previous step.

STAR entails the following novelties and computational advantages.

[1]In this framework, the maximals of the target intensity $I(x)$ are the intensities strictly greater than $I(x)$. If no maximals exist, the local reference white is set up as the target intensity.

First, differently from the sampling based Milano Retinex approaches, STAR does not define explicitly a set of sampled pixels, rather it picks up a set of intensity and distance values (i.e. the maximum intensity over each region and the mutual distance between regions), that in general do not correlate. In fact, the representative elements generally do not coincide with the pixels with maximum intensity over each segment. The number of intensity values sampled by STAR equals the number of segments, while the number of representative elements depends on the shapes of the segments. Both the numbers of the sampled intensities and of the representative elements are usually much lower than the number of pixels that need to be processed by the most sampling-based Milano Retinex algorithms. This procedure strongly decreases the number of operations to compute the lightness .

Second, the set of sampled information is independent of the target: differently from many Milano Retinex approaches, that need to repeat the sampling procedure at each target, STAR computes the spatial and color features relevant to the lightness only once, before the pixel-wise estimation of the lightness.

Third, the spatial terms weighting the maximals of the intensity target are pre-computed, cutting down the number of operations for the pixel-wise lightness computation.

The performance of STAR on image enhancement have been evaluated in two objective ways: first, by comparing some visual features, related to the understanding of the image content, on real-world color images before and after applying STAR; second, by studying if and how the color enhancement produced by STAR improves the accuracy of the well known Scale invariant feature transform (SIFT) algorithm [21] on the description and matching of pictures captured under difficult light conditions. These experiments show that STAR performs similarly to previous point-based Milano Retinex approaches [22], meaning that STAR is a good approximation of them, and that its color image enhancement effectively increases the SIFT accuracy. These tests also show that STAR, here equipped with the segmentation algorithm [23], is computationally more efficient than other point-based sampling Milano Retinex approaches. A comparison with some non-Milano Retinex enhancers is also proposed. Finally, the STAR performance has been judged by human observers, which looked at the images before and after STAR and expressed their preference in terms of details and content visibility.

The paper outlines as follows: Section II reports the state of the art on the sampling based Milano Retinex algorithms; Section III describes STAR; Section IV details the evaluation of the STAR performance; Section V draws the conclusions.

II. Related Work

This Section provides a brief survey of the sampling-based Milano Retinex algorithms, from which STAR is derived. More details can be found in [1] and [22].

Let I be an intensity channel of a RGB image. Let $S(I)$ indicate the support of I, i.e. the set of the spatial coordinates of the image pixels. Hereafter, I is represented as a function from $S(I)$ to the set of the intensity values, that here are normalized over $(0, 1]$, with zero excluded to prevent division

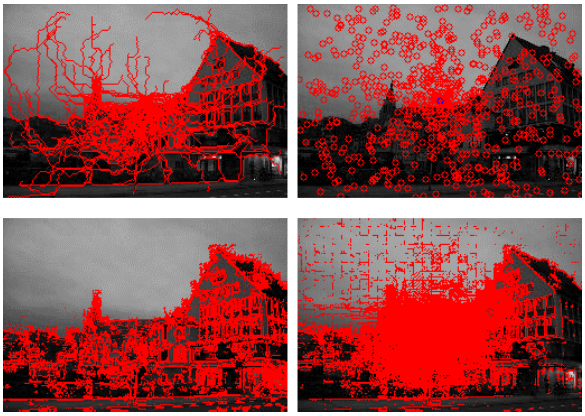

Fig. 1. In clockwise order: the sampling figures computed by ETR [11], RSR [13], GREAT [17] and GRASS [16] at the barycenter of the image support.

by zero. In particular, for any $y \in S(I)$ such that $I(y) = 0$, $I(y)$ is mapped on to $I(y)+\epsilon$, where ϵ is a positive, real value close to zero. This modification makes I strictly positive. Let L indicate the lightness of I.

The sampling based Milano Retinex approaches compute the lightness at any pixel x of $S(I)$ by the following equation:

$$L(x) = \frac{I(x)}{w(x)} \tag{1}$$

where $w(x)$ is the *local reference white* at x, and $w(x) > 0$. The value of $w(x)$ is obtained by processing the intensities of one or more pixels sampled around x. The union of these sets is the so-called *figure sampling* at x.

Equation (1) has been introduced by the Milano Retinex algorithm in [8], which explores the neighborhood of each target by n random paths $\gamma_1, \ldots, \gamma_n$ ending at x. Each path γ_k ($k = 1, \ldots, n$) is represented as a function from a discrete, ordered set $\Gamma_k := \{1, \ldots, l_k\} \subset \mathbf{N}$ to $S(I)$, with $\gamma_k(l_k) = x$, to $S(I)$. The lightness $L(x)$ is the average (over n) of the cumulative products of adjacent intensities along each path, with the threshold- and reset- constraints of the original Retinex theory:

$$L(x) = \frac{1}{n} \sum_{k=1}^{n} \prod_{t_k=1}^{l_k-1} \delta_k(R_{t_k}), \tag{2}$$

where $t_k \in \Gamma_k$, R_{t_k} is the ratio of the intensities $I(\gamma_k(t_{k+1}))$ and $I(\gamma_k(t_k))$, and $\delta_k : \mathbf{R}^+ \to \mathbf{R}^+$ is defined as:

$$\delta_k(R_{t_k}) = \begin{cases} R_{t_k} & \text{if } 0 < R_{t_k} \le 1-\varepsilon \\ 1 & \text{if } 1-\varepsilon < R_{t_k} \le 1+\varepsilon \\ R_{t_k} & \text{if } 1+\varepsilon \le R_{t_k} \le \dfrac{Q^{t_k}_{m_k}}{Q^{t_k-1}_{m_k}} \\ \dfrac{Q^{t_k}_{m_k}}{Q^{t_k-1}_{m_k}} & \text{if } R_{t_k} > \dfrac{Q^{t_k}_{m_k}}{Q^{t_k-1}_{m_k}} \end{cases} \tag{3}$$

Here, ε indicates a positive number close to zero, introduced to implement the threshold rule. The paper in [8] shows that, for the most real-world images, the threshold operation is unessential, thus it can be skipped. Within this approximation (i.e. $\varepsilon = 0$), Equation (2) reduces to Equation (1), where $w(x)$ is the average (over the number of paths) of the reciprocals of the maximum intensity over each path.

The Milano Retinex algorithms propose different sampling figures (see Figure 1). For instance, the algorithm [7] uses random paths starting from x and obeying the rules of a Brownian motion. The algorithms Termite Retinex (TR) [10], Energy-driven Termite Retinex (ETR) [11] and its fast version Light Energy-driven Termite Retinex (L-ETR) [12] introduce image-aware paths, i.e. curves that - differently from those proposed in [7], [8] - are not completely random, but depend on visual features of the image. Specifically, the TR, ETR and L-ETR paths are thought as walks of artificial termites, that move from x to other image places in search for the local reference white, and adhere as much as possible to image edges. Edges are key elements for lightness computation because of their relevant role in human color formation [2]. A mechanism inspired by the insect communication system spreads the swarm around each target in order to avoid the multiple exploration of a same image portion. TR and ETR differ to each other in the way to compute the paths (by local rules in TR, by minimizing an energy functional in ETR), while L-ETR is an optimized version of ETR. In [7], [10], [11], [12], the sampling figure is the union of the paths from x: the regions close to x are sampled more densely than those far away. This is in line with the empirical evidence that the color sensation depends more on the colors close to the target than on those located farther.

Based on this fact, Random Spray Retinex (RSR) [13] and its subsequent versions Spatially Weighted Random Spray Retinex (swRSR) [15], Light Random Spray Retinex [14], and Smart Light Random Spray Retinex [24] sample the neighborhood of x by random sprays, i.e. sets of pixels randomly sampled around x with radial density.

The algorithm GRASS [16] explores the target neighborhood by squares with random size and orientation and with the target located on the middle point of a square side. GRASS scans each square exhaustively by parallel segments, and detects over each segment the pixel maximizing the gradient magnitude and the closeness to the target. As in TR and ETR, the image scanning is optimized to avoid over-exploration. GRASS can be seen both as a path-based and a point based sampling method: in fact, the pixels selected from each square can be viewed as control points of a path, but, at the same time, since they are not connected, de facto they do not compose any curve.

In the method T-Rex [9], the sampling figure at any target is a set of pixels whose intensities, penalized by a term inversely proportional to their distance from the target, are greater than the target intensity.

The approach GREAT [17] proposes a global (i.e. not pixel-wise) computation of the sampling figure. The GREAT sampling figure is composed of edges whose gradient magnitude exceeds a pre-defined threshold. Locality is modeled in the lightness equation by taking into account the Euclidean distances between the target and the selected edges.

The equation of the local reference white $w(x)$ strongly depends on the sampling scheme.

The methods TR, ETR, L-ETR, RSR, swRSR and GRASS compute the term $\alpha(x) = \frac{1}{w(x)}$ in Equation (1) by this

formula:

$$\alpha^{\mathrm{TR,ETR,L-ETR,}}_{\mathrm{RSR,swRSR,GRASS}}(x) = \frac{\sum_{i=1}^{N} v_i(x) \frac{1}{\max\{I(y): y \in S_i\}}}{\sum_{i=1}^{N} v_i(x)} \quad (4)$$

where the S_i's, $i = 1, \ldots, N$, are the paths of TR, the random sprays (including also the target x) of RSR and swRSR, a stripe along the paths of ETR and L-ETR, and a set of pixels adjacent to the edges selected by GRASS. The v_i's are real, strictly positive numbers, introduced to tune the spatial interaction among colors. Specifically, in [8], TR, ETR, L-ETR, RSR, and GRASS the values of v_i's are constant (they are set to 1.0), while in swRSR they are inversely proportional to the distance of the target from the pixels with the maximum intensity over the spray. The algorithms [14] and [24] propose some optimizations of RSR, where the final value of the lightness is smoothed in order to reduce the chromatic noise due to the random sampling. The Brownian approach [7] uses the ratio-reset-threshold mechanism in Equations (2, 3), that, as observed before, on the most real-world images, is well approximated by Equation (1), and, precisely, by Equation (4) with the v_i's equal to 1.0.

The method T-Rex uses the following equation:

$$\alpha^{\mathrm{T-Rex}}(x) = \sum_{\bar{I}} \frac{1}{v_{\bar{I}}} \sum_{\bar{I}} v_{\bar{I}} \frac{1}{\bar{I}} \quad (5)$$

where $\bar{I}$ is an intensity value with $\bar{I} > v_{\bar{I}} I(x)$ and $v_{\bar{I}}$ is a weight, inversely proportional to the minimum distance between x and any pixel $y \in S$ such that $I = I(y)$. The sampling figure of T-Rex at x does not appear explicitly in this equation, but it is fundamental to detect the intensities of pixels satisfying the inequality $I > v_I I(x)$ and their weights. If no pixels satisfy this condition, then $w(x) = I(x)$.

The method GREAT implements the following formula:

$$\alpha^{\mathrm{GREAT}}(x) = \frac{\sum_{y \in S: I(y) > I(x)} v(x, y) \frac{1}{I(y)}}{\sum_{y \in S: I(y) > I(x)} v(x, y)} \quad (6)$$

where S the set the pixels adjacent to the sampled edges and to the target, and $v(x, y)$ is a spatial term weighting the contribution of the maximals of the target intensity.

All these Milano Retinex approaches contribute to investigate the concept of locality [4] and propose a large variety of color image enhancement techniques inspired by Retinex.

III. STAR

Let assume the notation introduced in Section II. STAR computes the lightness of I by two steps, called respectively *Global Processing* and *Pixel-wise Processing*. In the Global-Processing step, STAR segments the image and extracts spatial and color information from its segments; in the Pixel-wise Processing step, STAR processes the image pixel-wise according to the previous global analysis.

Global Processing

Let P be a segmentation of I, partitioning I into $M > 0$ regions $P_1, \ldots, P_M$. For each $j = 1, \ldots, M$, STAR computes the pair (I_j, L_j) where:

- I_j is the maximum value of the intensity over P_j, i.e.

$$I_j = \max\{I(u) : u \in P_j\}; \quad (7)$$

- L_j is the set of pixels of P_j maximizing the squared Euclidean distance transform over P_j, i.e. the elements of L_j maximize the Euclidean distance from the boundary ∂P_j:

$$L_j = \{u_j \in P_j : d(u_j, \partial P_j) = \max_{v \in \partial P_j} d(u_j, v)\} \quad (8)$$

where

$$d(u_j, v) = \frac{\| u_j - v \|^2}{\mathrm{diag}(I)^2}, \quad (9)$$

and $\mathrm{diag}(I)$ is the length of the diagonal of $S(I)$. The number of pixels of L_j depends on the geometry of L_j: if P_j is a circle or a square, then L_j contains only one element, that is the barycenter of P_j; if P_j is like a dumb-bell, then L_j contains two elements, located at its extremes. Due to the discrete nature of the data, in the STAR implementation, ∂P_j is the internal border of P_j, which is treated topologically as a closed set;

Given the set $\{(I_j, L_j)\}_{j=1,\ldots,M}$, STAR computes the map $\delta : \{1, \ldots, M\} \times \{1, \ldots, M\} \to [0, 1]$ measuring the minimum, mutual distances between the pixels of the sets L_j's, i.e.:

$$\delta(i, j) = \min\{d(u_i, v_j) : u_i \in P_i, v_j \in P_j\}. \quad (10)$$

Pixel-Wise Processing

For each pixel $x \in S(I)$, STAR detects the segment P_i such that $x \in P_i$ ($i \in \{1, \ldots, M\}$). Then it computes the subset $W_x \subset \Lambda = \{(I_j, L_j)\}_{j=1,\ldots,M}$ such that

$$W_x = \{(I_j, L_j) \in \Lambda : I_j > I(x)\}. \quad (11)$$

Finally, STAR computes the lightness $L(x)$ as follows:

$$L(x) = \frac{1}{\sum_{I_j \in W_x} (1 - \delta(i, j))} \sum_{I_j \in W_x} \frac{I(x)}{I_j}. \quad (12)$$

L is usually linearly rescaled to range over $\{0, \ldots, 255\}$.

Figure 2 illustrates the salient operations and the data of the global processing of STAR. The input image in (A) is partitioned in $M = 160$ regions, shown in (B). In (C), each segment is filled in with its maximum intensity value. The plot in (D) shows the global distance transform map Δ of the image segments, i.e. for each pixel p with spatial coordinates (r, c), $\Delta(r, c) = d(p, \partial P)$ where P is the segment including p. The red circles in (E) highlight the representative elements of the segments. This example clearly shows that elongated regions, like those at the bottom of the image, have more representative elements than convex segments, as for instance the triangular region on the top of the home at right. The image in (F) depicts the $M \times M$ matrix encoding the distances δ among the image segments.

The choices of the Euclidean distance, of the distance transform, and of the regional maximum intensities for the lightness computation are motivated as follows.

Use of the Euclidean distance: The use of the Euclidean

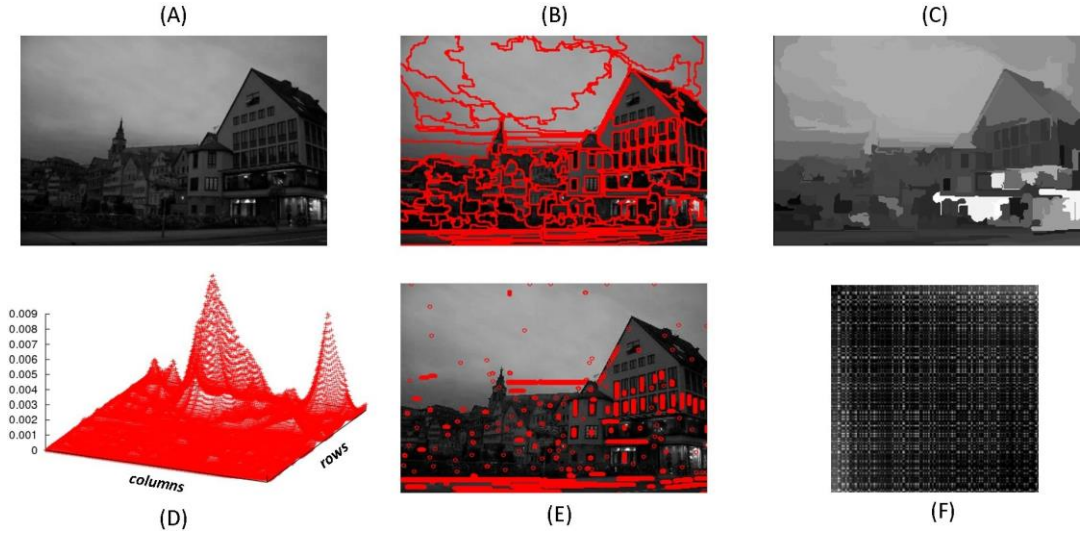

Fig. 2. (A) The red channel of an image from TEST35COLOR dataset; (B) its segmentation by [23] with $M = 160$; (C) image with the intensity values sampled by STAR: each segment of (B) is filled in with the gray level of the maximum value of the channel intensity over that segment (here, black = 0, white = 255); (D) the regional distance transforms Δ (see text); (F) the $M \times M$ map encoding the mutual distances between the segments.

distance to model locality is justified by many tests on human color perception [25], [4]. This metrics has been already employed in many previous Milano Retinex and Milano Retinex inspired approaches, like all the methods briefly described in Section II, the statistical approaches in [19] and [20], and the Retinex inspired color equalization approach in [26].

Use of the distance transform: Two different methods for computing approximated mutual distances between segments have been considered: they describe the position of each segment respectively by its barycenter and by the pixels maximizing the regional distance transform.

The barycenter of a region is commonly used to describe the position of the region in a compact and computationally efficient way. Nevertheless, measuring the distance between regions as the distance between their barycenters does not provide a sufficiently accurate estimate of the spatial information required by Milano Retinex. In fact, while the representative elements of STAR are always inside the region, the barycenters may also fall out of them. This is the case of a non-convex region, let's say P_i, whose barycenter falls out of P_i, and coincides with the barycenter of another region, let's say P_j, adjacent to P_i ($i = j$). The distance between the barycenters is null: this information poorly describes the spatial relative position of P_i and P_j and cannot be used to weight the maximum intensities of P_i and P_j. This example clearly shows that pixels inside the segment describe better the location of the segment. Among the many possible choices of internal pixels (e.g. points randomly sampled inside the region), STAR considers the pixels maximizing the distance transform over the segment, i.e. the most internal points.

Use of the Maximals of the Target Intensity: The computation of the lightness as weighted average of the maximals of the target intensity is inspired by previous Milano Retinex approaches, like GREAT and T-Rex (see Section II).

STAR shares with GREAT also the extraction of spatial and color information *independently* of the target. This procedure allows to cut down the computational cost due to the *pixel-wise* sampling procedure. Nevertheless, as already highlighted in Section I, STAR strongly differs from GREAT and from the approaches mentioned in Section II, because it does not define explicitly a sampling figure, rather it samples a set of relevant intensity values and a set of pixels, not necessarily correlated to the sampled intensities and employed to estimate the spatial weights of the maximals in the lightness equation. This characteristic makes STAR an innovative spatial color algorithm inspired by Milano Retinex.

The current implementation of STAR embeds the segmentation algorithm [23]. For each color channel, this algorithm represents $S(I)$ by a planar weighted graph, where the nodes are the image pixels, the edges connect pixels which are neighbors in the 8-connected sense, and the weights on the edges measure the absolute difference of the intensities of the connected pixels. The segmentation divides $S(I)$ in M regions $P_1, \ldots, P_M$. For any $i, j = 1, \ldots, M$, $i = j$ let D_i indicate the largest weight in the minimum spanning tree of P_i, and let D_{ij} the minimum weight of the edges connecting P_i with P_j. The regions P_i and P_j are characterized as follows: the value of D_{ij} is large *relative* to the values of D_i and D_j with respect to a threshold depending on the size of the segments and tunable by the user. This segmentation has been choosen among others, e.g. [27], [28] for three main reasons: it exploits local spatial intensity variations, that are key-points in Retinex theory; it captures perceptual important regions; finally, it has a short execution time, thanks to its complexity nearly linear with $|I|$. Low complexity and short runtime are very important features, since STAR aims at reducing the computational burden and the execution time of point-based sampling Milano Retinex approaches. Within this segmentation, the computational complexity of STAR is $O(M|I|)$, and since for the most real-world images $M \quad |I|$,

then $O(M|I|) \quad O(|I|^2)$ (see also Section IV-E).

IV. Experiments

Quantifying the performance of an image enhancement algorithm is a difficult task, mainly because of the lack of agreed measures [4], [29]. An *universally optimal* image enhancer does not exist: in fact, any image enhancer is usually designed for specific tasks [30], therefore its performance must be estimated according to the application for which the enhancer has been developed. Consequently, different evaluation measures have been proposed in the literature, e.g. [31], [32], [33], [34].

In this work, the performance of STAR have been evaluated by three ways: first, by measuring and comparing a set of numerical image features before and after applying STAR (Subsection IV-A); second, by measuring the benefit of the image enhancement produced by STAR on the unsupervised description and matching of images performed by SIFT [21] (Subsection IV-B); finally, by perceptual tests with human observers that compare the images before and after applying STAR and state which image has the most *understandable* content and the highest detail visibility (Subsection IV-C).

The performance analysis has been carried out on two datasets with real-world color images (Subsection IV-D). The results are discussed in Subsection IV-E.

A. Objective Visual Feature Comparison

The following numerical image features are usually modified by enhancers: the mean brightness (f_0), the multi-resolution contrast [35] (f_1), and the distance of the probability density function of the image brightness from the flatness (f_2). These features are related to the human understanding of the image content, and they have been already employed to evaluate previous Milano Retinexes, e.g. [11], [16], [17], [9].

The f_i's are computed on the brightness B of any color image J: here B is the mono-chromatic image obtained by averaging pixel-by-pixel the channel intensities of J. Precisely, f_0 is the mean value of B. The multi-resolution contrast f_1 is the mean value of the local contrast over a pyramid of images, built up from a sequential down-scaling of B by 0.5 as proposed in [35]. The value of f_2 is the L^1 distance between the intensity probability density function of B and of a variable uniformly distributed over (0, 1]. It measures the *flatness* of the histogram of B.

The values of f_0 and f_1 (f_2, resp.) of the color lightness are expected to be higher (lower, resp.) than those of the input image. The amount of the variations of the f_i's depends on the input image, on the sampling scheme and on the lightness equation. In particular, if the image content is already clear (resp. unintelligible), then the modification due to STAR will be slight (resp. very evident), and the variation of the f_i's before and after STAR will be negligible (resp. remarkable). An example is given in Figure 3.

The perceptual differences between the input and the output images are measured here by the CIELab distance ΔE between any RGB image and its color lightness [36], assuming a D65 illuminant. The value of ΔE is the mean value of the L^1 distances between the color coordinates of these images expressed in the Lab space. The higher ΔE, the higher the amount of perceptual changes between the observed images is.

STAR performance is also evaluated through the measures NIQE [33] and BRISQUE [34][2]. NIQE and BRISQUE quantify image distorsions that may be introduced by image processoing algorithms and that may affect the *naturalness* of the image. In this framework, image naturalness is defined as the regularity of low order image statistics, computed globally or locally. Regularity is modeled by multivariate Gaussians of the brightness estimated from training sets of images, that in [34] also include information about subjective judgements. NIQE and BRISQUE measure departures from this regularity: algorithms which decrease NIQE and BRISQUE, increase the image naturalness.

B. SIFT Description and Matching before and after STAR

The algorithm SIFT (Scale Invariant Feature Transform) [21] describes the visual content of an image by a set of histograms encoding the intensity distribution around a set of high-contrast pixels, termed *key-points* and invariant to image re-scaling. SIFT takes as input a gray level image I, and extracts its key-points as the local extrema of a difference of Gaussian functions computed over a series of re-scaled and smoothed versions of I. Key-points are filtered to retain only high-contrast pixels and edges with high gradient magnitude. SIFT associates each key-point to an histogram with 128 entries, representing the intensity distribution in a 16×16 key-point neighbourhood. SIFT matches two images I and I' by individually comparing their descriptors through the Euclidean distance. The best candidate match for each keypoint of I is its nearest neighbor among the key-points of I'. A candidate match is retained only if the ratio of its distance from the closest neighbor to the distance of the second closest is smaller than a fixed threshold τ (here $\tau = 0.7$).

Since STAR works as an image enhancer, and, in particular, it improves the visibility of the image details and thus its contrast, it can be very useful to detect keypoints in images captured under difficult illuminant conditions, for instance due to dark environments, backlight, low exposure time of the camera. This paper studies how the SIFT performance, measured in terms of numbers of key-points and matches changes when such images are enhanced by STAR.

C. Subjective Image Comparison

Objective measures, like the f_i's, provide numerical estimates of the variations of visual perceptual features, but they are not able to model the complex mechanism of the human perception, which involves both physical and psychological cues. For this reason, subjective evaluation is often recommended to assess the performance of an image enhancers in applications like image visualization and video restoration, e.g. [37], [38]. Therefore, this work reports a subjective analysis

[2]NIQE and BRISQUE have been computed by the routines implemented in MatLab Release 2017b.

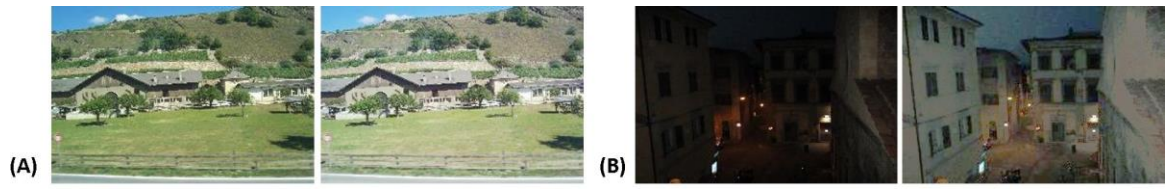

Fig. 3. From TEST35COLOR: two images and their color lightness computed by STAR. In (B) the input image (on left) is poorly readable and the image enhancement effects of STAR (on left) are much more evident than those in (A), where the input image is already clear.

of the STAR performance on a set of natural images (see Section IV-D). Thirty volunteers, with different gender, age and profession, were asked individually to look at each pair of images and to select the image with the most readable content and most visible details. Each image was displayed on a screen along with its version processed by STAR. The order of the images on the screen has been randomly choosen at the begin of the experiments and it was the same for all the observers.

D. Datasets

The database TEST35COLOR is composed of 35 real-world color images, acquired by different commercial cameras and depicting indoor and outdoor environments (e.g. Figure 3). Despite their small size, ranging from 202×149 to 260×252 pixels, these images are suitable to test spatial color algorithms, because of the different distributions and size of bright and dark areas depicted in. Moveover, they have already been employed to test other Milano Retinexes, e.g. [11], [12].

The dataset FLIR20 is composed of two sets of 20 real-world color pictures, hereafter denoted by FLIR20-Le and FLIR20-He. These images, with size 640×480 pixels, have been captured by the FLIR camera FireFly FMVU-03MTC-CS [3] with automatic exposure control disabled. The images of FLIR20-Le and FLIR20-He depict the same scene, but they differ in the content visibility. In fact, the images of FLIR20-Le have been captured with low exposure, so that they appear very dark and their content is poorly visible. The images of FLIR20-He have been acquired with a higher exposure and their content is thus more understandable (see Figure 4). The suffix Le (He, resp.) in the database name means just "Low-Exposure" ("High Exposure", resp.).

The measures in Section IV-A have been evaluated on both databases. Perceptual tests have been carried on TEST35COLOR, which includes a larger variety of images than FLIR20. Finally, the color enhancement of STAR as a pre-processing step of SIFT to match images taken under difficult light conditions has been evaluated on FLIR20.

E. Results and Discussion

Tables I, II and III reports the objective measures and the evaluation of STAR in terms of SIFT performance, along with a comparative analysis of STAR against Milano Retinex and non-Milano Retinex approaches. The comparison with Milano Retinexes aims at proving that STAR effectively approximates point-based sampling approaches, while that with non-Milano Retinexes aims at giving a broad picture of STAR in the image enhancement context. To enable a qualitative analysis of the results, the images from TEST35COLOR and FLIR20 displayed in the figures from 1 to 7 are available along with their enhanced versions at full resolution in the supplementary material attached to this work. For FLIR20, the figures with overlaid SIFT key-points are also provided.

The Milano Retinex methods matched against STAR are GREAT, GRASS, RSR and L-ETR. Here, the parameters of GREAT, GRASS and L-ETR have been set as suggested in [17], [16], [12] respectively, while the number N of sprays and the number m of pixels per spray in RSR are specified in the tables. The experiments on TEST35COLOR considers all these methods, while that on FLIR20 excludes GRASS and L-ETR because of their long execution time, which is not comparable with that of STAR.

The non-Milano Retinex methods considered here are: histogram equalization (HE), contrast-adaptive histogram equalization (CLAHE), naturalness preserved enhancement algorithm for non-uniform illumination images (NPEA) [39], non uniform low-light image enhancement (LIME) [40] and non local retinex contrast enhancement (NL-RETINEX) [41].
HE and CLAHE improve contrast by manipulating the distributions of the image brightness (HE-L, CLAHE-L) or of the image color channels (HE-C, CLAHE-C). Specifically, HE spreads globally the intensity distributions, while CLAHE clips the input histograms at a predefined threshold X and locally re-distributes the image intensities. NPEA, LIME and NL-RETINEX [4] address the (ill-posed) problem of reflectance-illumination image decomposition in a constrained solution space and enhance the input image by lowering or removing its illuminant component. In particular, NPEA imposes the preservation of the local relative order of the illumination values. LIME implements a variational model based on the dark channel prior widely employed for dehazing [42]. Inspired by the work [43], which hypothesizes smoothly varying illuminants, NL-RETINEX achieve image enhancement by using non local (NL) differential operators.

Tables I and II show that STAR modifies the input image by increasing its brightness, contrast and histogram flatness, and this behaviour is reported also by the other algorithms compared with STAR. The exact amount of the f_i's depends on the way the image colors are processed and thus varies from

[3]https://www.ptgrey.com/firefly-mv-03mp-color-usb-20-micron-mt9v022

[4]NPEA, LIME and NL-RETINEX codes and parameters used here are available respectively at https://shuhangwang.wordpress.com/2015/12/14, https://sites.google.com/view/xjguo/lime and https://it.mathworks.com/matlabcentral/fileexchange/47562-non-local-retinex?focused=3833114&tab=function (without downsampling).

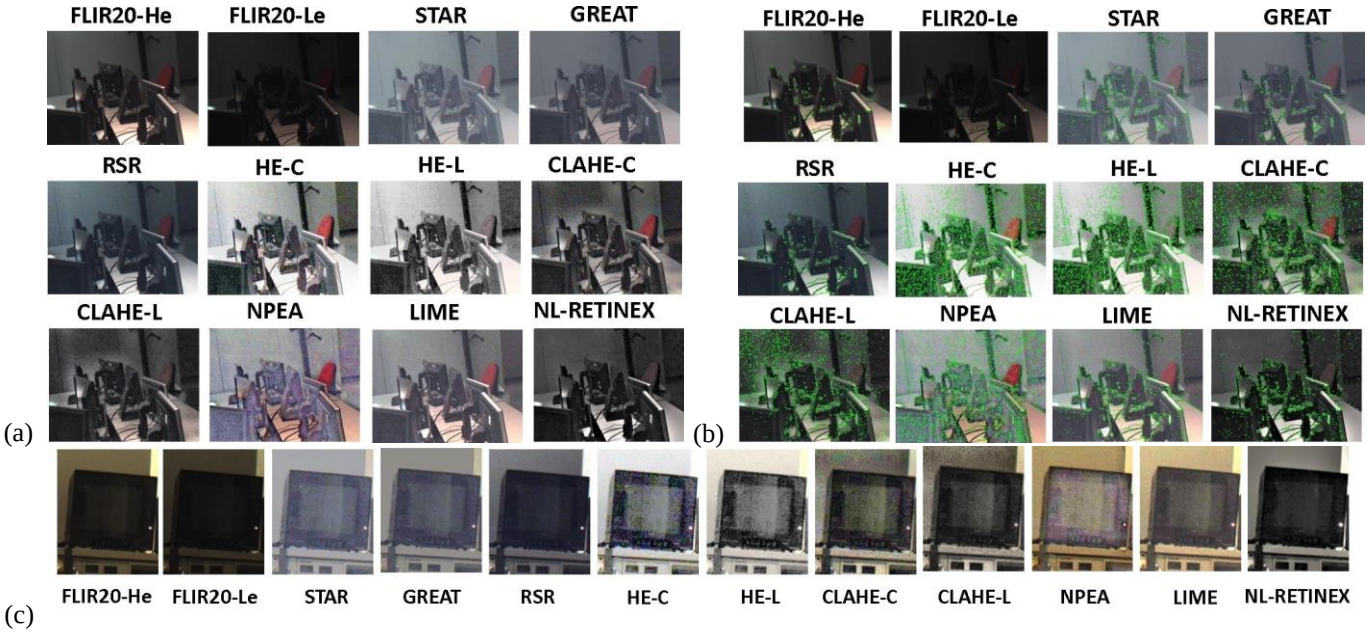

Fig. 4. From FLIR20: (a) an image from FLIR20-LE, its under-exposed version from FLIR20-HE and its enhanced versions by STAR and other algorithms; (b) SIFT key-points (in green) detected on the images in (a); (c) the portion of an image from FLIR20 and of its enhanced versions. See also the Supplementary Material attached to this paper.

Fig. 5. From TEST35COLOR: An image and its STAR color lightnessess obtained by varying the parameters of the segmentation [23].

algorithm to algorithm. Examples of image enhancement by STAR and/or by other enhancers are depicted in Figures 3, 4, 5, 6, 7 (see also Supplementary material).

On TEST35COLOR, the parameters of the segmentation algorithm (i.e. three values controlling the size of the segments and the image smoothness) have been tuned to produce a different number of segments (see Figure 5). The objective measures of STAR obtained by varying M are shown in Table I(A), where M denotes the mean number of image segments. On these data, a finer segmentation produces on average higher (smaller, resp.) values of f_0 and f_1 (f_2, resp.) than a coarser one. It is to note that STAR performs image enhancement whatever the segmentation algorithm is. Anyway, different segmentations may yield different enhancement levels. There are two limit cases: (*i*) $M = 0$: STAR behaves like scale-by-max; (*ii*) $M = |I|$: STAR does not have any computational advantage, since its computational complexity is $\mathcal{O}(|I|^2)$. A general recommendation is to use a not too coarse and not too fine segmentation. The effects of the segmentation granularity (i.e. number of segments) on L can be inferred from Equation (12). These effects are related to the shape of the segments as well. Precisely, a fine segmentation tends to split any almost uniform region R more than a coarse segmentation. This means that, for any target x, more intensity values close to each other, let's say $I_1, \ldots, I_K$, are sampled from R and contribute to $L(x)$. Within a coarse segmentation, where R is a unique segment, only the intensity $I_{max} = \max(I_1, \ldots, I_K)$ contributes to $L(x)$. The different values of $L(x)$ output by these two segmentations depend on the spatial terms weighting the I_j's and I_{max}, i.e. on the shape of R and more generally on the image content. From Equation (12) it follows that, if the spatial terms weighting the I_j's and I_{max} in the finer and coarser segmentations are comparable, then the finer segmentation outputs the highest value of $L(x)$. In other cases, the result depends on the values of the I_j's and I_{max} and of their distances from x.

Tables I and II show that STAR performs similarly to GRASS, RSR and in particular to GREAT, meaning that STAR is effectively a good approximation of these point-based sampling approaches. L-ETR performs very differently from STAR, RSR, GRASS and GREAT because of its different sampling mechanism (path-based versus point-based). Among the point-based sampling methods, STAR yields the highest values of f_0 and ΔE and the smallest value of f_1, while the value of f_2 is close to those of the other point-based sampling approaches (here we assumed $M = 160$ for TEST35COLOR). Among the non-Milano Retinex algorithms, NPEA is that reporting more similar values to STAR.

On average, the effects of the color enhancements on the image regularity (in terms of NIQE and BRISQUE) differ on

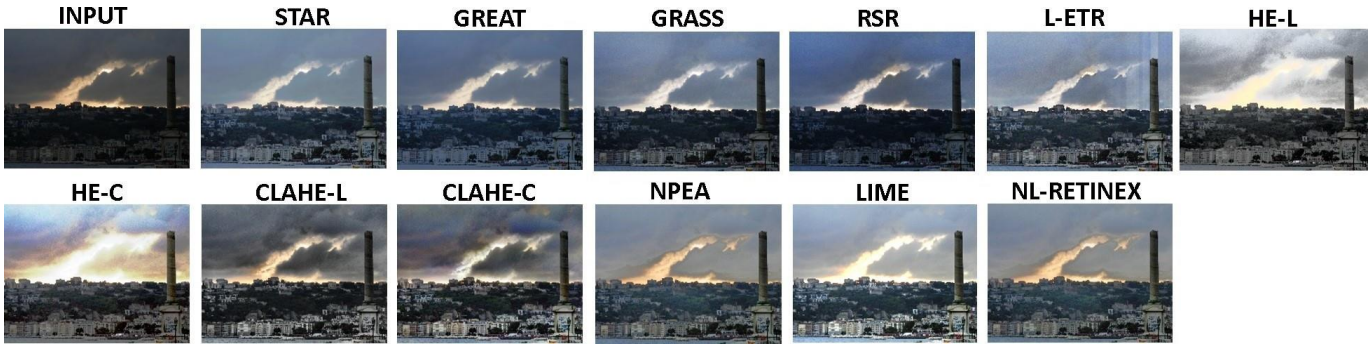

Fig. 6. TEST35COLOR: an example of color lightness computed by STAR in comparison with other approaches

the two datasets. On TEST35COLOR, all the enhancers except GREAT decrease the naturalness of the input images: LIME provides the worst results, while STAR increase the NIQE and BRISQUE values of the input images by the 1.82% and 0.57% respectively. On FLIR20, all the algorithms improve the image naturalness. On these data, STAR decreases the NIQE and BRISQUE values of the input images respectively by the 7.96% and by the 19.86%. CLAHE-L and RSR achieve the lowest values of NIQE and BRISQUE respectively.

Despite the objective measures reported in Tables I and II have a similar trend, the enhanced images exhibit important visual differences. A qualitative analysis shows that the non-Milano Retinex methods often over-enhance the images highliting unessential details and noise (see Supplementary Material) much more than STAR and Milano Retinexes. Artifacts introduced by over-enhancing are particularly visible in the images processed by HE and CLAHE (see Figure 4(c)). The presence of noisy edges may adversely affect the image description and matching, providing a verbose and often meaningless characterization of the image content, as proved by the experiments with SIFT. In this framework, the images of FLIR20-He and their enhanced versions have been compared with those of FLIR20-Le by SIFT. For this task, all the images have been converted to gray-scale and input to SIFT for description and comparison. Table III reports the mean number of key-points (N_{kp}), the mean number of retrieved matches (N_{tm}) and the mean number of *correct* matches (N_{cm}), where a match is *correct* if it links corresponding key-points. On average, the value of N_{kp} on FLIR20-Le is much smaller than that on FLIR20-He. For some images of FLIR20-He, no key-points have been found (e.g. Figure 7(A)). All the enhancers considered here report a very high value of N_{kp}, often much greater than that of FLIR20-He. This is because these algorithms not only enhance actual, important details, but also the noisy pixels due to the camera low-exposure. On FLIR20, the worst performance is obtained by the non-Milano Retinex approaches (e.g. Figures 4(A) and (B), and Supplementary Material), that generally return a high percentage of noisy key-points and false matches. Among the Milano Retinex algorithms, the performance of STAR are once again very close to that of GREAT. Anyway, all the enhancers considered here improve the SIFT performance, retrieving at least 3 matches between corresponding images.

The subjective analysis reports that the 59.71% of the observers prefers the images of TEST35COLORS processed by STAR, the 24.01% selected the input images, while the 16.28% had no preferences. In general, when the brightness and the contrast are low and the intensity distributions are peaked, people prefer the enhanced image. When the content of the input image is already readable and thus it differs from the output at most for the presence of a color dominant (which is partially removed by STAR), then observers have no preferences or select one image upon their aesthetic feeling. For instance, in Figure 3, in (A) the 74.7% of the observers preferred the input image, while in (B) the 83.3% preferred the color lightness.

TABLE I
OBJECTIVE EVALUATION ON (A) TEST35COLOR AND (B) FLIR20.

(A)

Algorithm	f_0	f_1	$f_2[\times 10^{-3}]$	ΔE
INPUT _	64.30	15.34	4.13	-
STAR ($\overline{M} = 160$)	98.62	19.98	3.16	20.82
STAR ($\overline{M} = 137$)	94.87	19.83	3.19	18.32
STAR ($\overline{M} = 69$)	91.76	19.49	3.27	17.04
GREAT	96.93	19.74	3.29	18.88
GRASS	89.19	20.81	3.14	16.65
RSR ($N = 20$, $m = 150$)	87.41	20.66	3.20	14.33
L-ETR	104.30	23.13	2.67	25.15
HE-C	131.62	25.65	1.55	43.99
HE-L	129.96	26.04	2.04	38.94
CLAHE-C ($X = 4$)	100.72	27.09	2.23	27.86
CLAHE-L ($X = 4$)	93.61	26.31	2.45	23.49
NPEA	99.60	20.49	3.04	24.56
LIME	125.81	27.81	2.21	32.92
NL-RETINEX	84.18	22.32	3.08	23.34

(B)

Algorithm	f_0	f_1	$f_2[\times 10^{-3}]$	ΔE
FLIR20-Le (none)	32.39	6.67	6.00	-
FLIR20-He (none)	73.65	15.51	4.05	28.28
STAR ($\overline{M} = 1538$)	124.54	12.43	4.73	52.13
GREAT	106.71	12.59	4.71	43.18
RSR ($N = 20$, $m = 200$)	87.00	15.09	4.06	34.93
HE-C	139.14	24.65	3.95	57.45
HE-L	138.37	24.76	4.72	54.89
CLAHE-C ($X = 8$)	78.20	20.38	3.35	32.02
CLAHE-L ($X = 8$)	69.98	18.71	3.73	25.97
NPEA	118.84	15.89	4.26	51.50
LIME	113.18	17.18	3.75	47.51
NL-RETINEX	68.56	18.90	3.57	28.29

From the computational view-point, STAR (with $O_{STAR} = O(M|I|)$) is the most efficient algorithm among L-ETR,

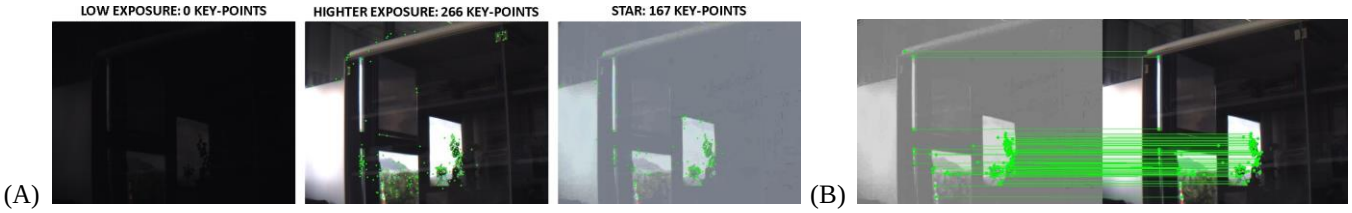

Fig. 7. FLIR20: (A) SIFT key-points of an image from FLIR20-Le (left), of its corresponding version in FLIR20-He (middle) and of its color lightness by STAR (right); (B) key-point matching by SIFT between the image from FLIR20-Le enhanced by STAR and its version from FLIR20-He.

TABLE II
NIQE AND BRISQUE ON TEST35COLORS (A) AND FLIR20 (B).

	Algorithm	NIQE	BRISQUE
(A)	INPUT	7.01	26.28
	STAR ($\overline{M} = 160$)	7.14	26.43
	GREAT	7.19	24.87
	GRASS	6.82	22.59
	RSR (N = 20, m = 150)	7.16	23.63
	L-ETR	7.24	24.29
	HE-C	7.79	26.91
	HE-L	7.81	28.51
	CLAHE-C (X = 4)	7.81	27.68
	CLAHE-L (X = 4)	7.79	26.91
	NPEA	7.33	28.36
	LIME	8.26	30.24
	NL-RETINEX	7.49	26.86
	Algorithm	NIQE	BRISQUE
(B)	FLIR20-He (none)	3.79	39.23
	FLIR20-Le (none)	4.65	48.45
	STAR	4.28	38.93
	GREAT	4.23	41.12
	RSR (N= 20, m = 200)	4.46	25.93
	HE-C	4.48	40.92
	HE-L	4.45	40.89
	CLAHE-C (X = 8)	3.85	35.67
	CLAHE-L (X = 8)	3.78	38.12
	NPEA	4.46	42.15
	LIME	4.98	46.01
	NL-RETINEX	4.17	32.37

TABLE III
FLIR20: EVALUATION AND COMPARISON OF STAR COLOR ENHANCEMENT IN TERMS OF SIFT PERFORMANCE.

Dataset	N_{kp}	N_{tm}	N_{cm}	N_{cm}/N_{tm})
FLIR20-He	588	–	–	–
FLIR20-Le	117	68	41	0.60
STAR	680	154	140	0.91
GREAT	625	166	152	0.92
RSR (N = 20, m = 200)	807	232	214	0.92
HE-C	2660	385	165	0.43
HE-L	2619	369	155	0.42
CLAHE-C	2227	436	201	0.46
CLAHE-L (X = 8)	1670	416	197	0.47
NPEA	2233	391	113	0.29
LIME	1260	444	270	0.61
NL-RETINEX	1367	364	161	0.44

GRASS, RSR and GREAT. L-ETR has the highest computational complexity, i.e. $O_{L-ETR} = O(|I|^2 \log |I|)$. GRASS performs computationally better: in the worst case, i.e. when all the sampled squares have size $|I|$, the complexity of GRASS is $O_{GRASS} = O(|I|^2)$, but in general, $O_{GRASS} \leq O_{L-ETR}$. The computational complexities of RSR and GREAT are much lower than O_{L-ETR} and O_{GRASS}: they are respectively $O_{RSR} = O(Nm|I|)$ and $O_{GREAT} = O(\Sigma|I|)$, with Σ being the size of the GREAT sampling figure. In the most real-world applications, both Nm and Σ are greater than M, so that O_{STAR} is lower than O_{RSR} and O_{GREAT}. In particular, on a standard PC with CPU Intel(R) Xeon(R) CPU E3-1245 v6 @ 3.70GHz, on the images of FLIR20, the current implementations of L-ETR and GRASS require more than one hour, versus about 20 minutes of GREAT (with $\Sigma \approx \frac{3}{10}|I|$) and RSR (with N = 20, m = 200), and versus 11 seconds of STAR (with M = 1538). In these experiments, segmentation requires on average less than 0.1 seconds per channel.

Finally, Figure 8 emphasizes an advantage of the deterministic sampling of STAR against the random sampling of RSR, i.e. the absence of chromatic noise in the image processed by STAR.

Fig. 8. (a, left) A toy image displaying a bright square over a dark one, and its color lightnesses by (b, left) STAR and (c, left) RSR (N = 20, m = 100). The small squares in (a, b, c, right) are enlargements of the bottom right corner of the images on left. Enlargements have been equalized to allow the visualization of the noise generated by RSR. This noise is totally absent in the image processed by STAR.

V. CONCLUSIONS

This paper presented STAR, a spatial color algorithm inspired by Retinex principles. STAR introduces a novel scheme for the lightness computation based on non-correlated spatial and color information extracted from clusters of pixels instead of single pixels. The experiments reported here show that STAR is a good approximation of point-based sampling Milano Retinex approaches, in particular of GREAT, and that its color enhancement may be successfully used to detect visual features relevant to describe and match images captured under difficult light conditions. Future work will investigate other applications of STAR to computer vision, e.g. object recognition, image retrieval, image/video restoration, and it will include a further optimization of the STAR code.

ACKNOWLEDGMENTS

The author would like to thank Stefano Messelodi, Paola Lecca, Alessio Xompero for their help on image segmentation, algorithm evaluation and comparison.

REFERENCES

[1] A. Rizzi and C. Bonanomi, "Milano Retinex family," *Journal of Electronic Imaging*, vol. 26, no. 3, pp. 031 207–031 207, 2017.

[2] E. H. Land, John, and J. McCann, "Lightness and Retinex theory," *Journal of the Optical Society of America*, vol. 1, pp. 1–11, 1971.

[3] J. J. McCann, "Retinex algorithms: Many spatial processes used to solve many different problems," *Electronic Imaging*, vol. 2016, no. 6, pp. 1–10, 2016.

[4] J. McCann and A. Rizzi, *The Art and Science of HDR Imaging*. John Wiley, 2011.

[5] M. A. Wallach, "On psychological similarity." *Psychological Review*, vol. 65, no. 2, p. 103, 1958.

[6] B. Funt, F. Ciurea, and J. McCann, "Retinex in MATLAB," *Journal of electronic imaging*, vol. 13, no. 1, pp. 48–57, 2004.

[7] D. Marini and A. Rizzi, "A computational approach to color adaptation effects," *Image and Vision Computing*, vol. 18, no. 13, pp. 1005 – 1014, 2000.

[8] E. Provenzi, L. D. Carli, A. Rizzi, and D. Marini, "Mathematical definition and analysis of the Retinex algorithm," *Journal of the Optical Society of America A: Optics and Image Science, and Vision*, vol. 22, no. 12, pp. 2613–2621, 2005.

[9] M. Lecca, C. M. Modena, and A. Rizzi, "T-Rex: A Milano Retinex Implementation based on Intensity Thresholding," in *Int. Workshop on Computational Color Imaging*. Springer, Cham, 2017, pp. 68–79.

[10] G. Simone, G. Audino, I. Farup, F. Albregtsen, and A. Rizzi, "Termite Retinex: a new implementation based on a colony of intelligent agents," *Journal of Electronic Imaging*, vol. 23, no. 1, 2014.

[11] M. Lecca, A. Rizzi, and G. Gianini, "Energy-driven path search for Termite Retinex," *Journal of Optical Society of America, A*, vol. 33, p. 1, 2016.

[12] G. Simone, R. Cordone, R. P. Serapioni, and M. Lecca, "On edge-aware path-based color spatial sampling for Retinex: from Termite Retinex to Light Energy-driven Termite Retinex," *Journal of Electronic Imaging*, vol. 26, no. 3, pp. 031 203–031 203, 2017.

[13] 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," *Trans. Img. Proc.*, vol. 16, no. 1, pp. 162–171, Jan. 2007.

[14] Banić, N. and Lončarić, S., "Light Random Sprays Retinex: exploiting the noisy illumination estimation," *IEEE Signal Processing Letters*, vol. 20, no. 12, pp. 1240–1243, 2013.

[15] M. Lecca and A. Rizzi, "Tuning the locality of filtering with a spatially weighted implementation of Random Spray Retinex," *JOSA A*, vol. 32, no. 10, pp. 1876–1887, 2015.

[16] M. Lecca, A. Rizzi, and R. P. Serapioni, "GRASS: A gradient-based random sampling scheme for Milano Retinex," *IEEE Transactions on Image Processing*, vol. 26, no. 6, pp. 2767–2780, 2017.

[17] M. Lecca, A. Rizzi, and R.P. Serapioni, "GREAT: a gradient-based color-sampling scheme for Retinex," *J. Opt. Soc. Am. A*, vol. 34, no. 4, pp. 513–522, 2017.

[18] G. Gianini, A. Rizzi, and E. Damiani, "A Retinex model based on absorbing Markov chains," *Information Sciences*, vol. 327, pp. 149–174, 2016.

[19] G. Gianini, A. Manenti, and A. Rizzi, "QBRIX: a quantile-based approach to Retinex," *J. Opt. Soc. Am. A*, vol. 31, no. 12, pp. 2663–2673, Dec 2014.

[20] G. Gianini, M. Lecca, and A. Rizzi, "A population-based approach to point-sampling spatial color algorithms," *JOSA A*, vol. 33, no. 12, pp. 2396–2413, 2016.

[21] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *Int. Journal of Computer Vision*, vol. 60, no. 2, pp. 91–110, 2004.

[22] M. Lecca, G. Simone, C. Bonanomi, and A. Rizzi, "Point-based spatial colour sampling in milano-retinex: a survey," *IET Image Processing*, vol. 12, pp. 833–849(16), June 2018.

[23] P. F. Felzenszwalb and D. P. Huttenlocher, "Efficient graph-based image segmentation," *Int. journal of computer vision*, vol. 59, no. 2, pp. 167–181, 2004.

[24] Banić, N. and Lončarić, S., "Smart light random memory sprays Retinex: a fast Retinex implementation for high-quality brightness adjustment and color correction," *JOSA A*, vol. 32, no. 11, pp. 2136–2147, 2015.

[25] O. Creutzfeldt, B. Lange-Malecki, and K. Wortmann, "Darkness induction, Retinex and cooperative mechanisms in vision," *Experimental Brain Research*, vol. 67, no. 2, pp. 270–283, 1987.

[26] A. Rizzi, C. Gatta, and D. Marini, "A new algorithm for unsupervised global and local color correction," *Pattern Recogn. Lett.*, vol. 24, no. 11, pp. 1663–1677, Jul. 2003.

[27] M. Sonka, V. Hlavac, and R. Boyle, *Image processing, analysis, and machine vision*. Cengage Learning, 2014.

[28] R. Zwiggelaar and E. R. Denton, "Texture based segmentation," in *Int. Workshop on Digital Mammography*. Springer, 2006, pp. 433–440.

[29] M. Pedersen and J. Y. Hardeberg, "Full-reference image quality metrics: Classification and evaluation," *Foundations and Trends® in Computer Graphics and Vision*, vol. 7, no. 1, pp. 1–80, 2012.

[30] M. E. Celebi, M. Lecca, and B. Smolka, *Color Image and Video Enhancement*. Springer, vol. 4.

[31] Y. A. Al-Najjar, D. C. Soong *et al.*, "Comparison of image quality assessment: PSNR, HVS, SSIM, UIQI," *Int. Journal of Scientific & Engineering Research*, vol. 3, no. 8, p. 1, 2012.

[32] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE transactions on image processing*, vol. 13, no. 4, pp. 600–612, 2004.

[33] A. Mittal, R. Soundararajan, and A. C. Bovik, "Making a completely blind image quality analyzer," *IEEE Signal Processing Letters*, vol. 20, no. 3, pp. 209–212, March 2013.

[34] A. Mittal, A. K. Moorthy, and A. C. Bovik, "No-reference image quality assessment in the spatial domain," *IEEE Transactions on Image Processing*, vol. 21, no. 12, pp. 4695–4708, Dec 2012.

[35] A. Rizzi, T. Algeri, G. Medeghini, and D. Marini, "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*, Aachen, 2004, pp. 187–192.

[36] C. Oleari, *Standard Colorimetry: Definitions, Algorithms and Software*. John Wiley & Sons, 2016.

[37] A. Rizzi and M. Chambah, "Perceptual color film restoration," *SMPTE Motion Imaging Journal*, vol. 119, no. 8, pp. 33–41, 2010.

[38] A. Rizzi, A. J. Berolo, C. Bonanomi, and D. Gadia, "Unsupervised digital movie restoration with spatial models of color," *Multimedia Tools and Applications*, vol. 75, no. 7, pp. 3747–3765, 2016.

[39] S. Wang, J. Zheng, H.-M. Hu, and B. Li, "Naturalness preserved enhancement algorithm for non-uniform illumination images," *IEEE Transactions on Image Processing*, vol. 22, no. 9, pp. 3538–3548, 2013.

[40] X. Guo, Y. Li, and H. Ling, "Lime: Low-light image enhancement via illumination map estimation," *IEEE Transactions on Image Processing*, vol. 26, no. 2, pp. 982–993, 2017.

[41] D. Zosso, G. Tran, and S. J. Osher, "Non-local Retinex - a unifying framework and beyond," *SIAM Journal on Imaging Science*, vol. 8, no. 2, pp. 787–826, 2015.

[42] K. He, J. Sun, and X. Tang, "Single image haze removal using dark channel prior," *IEEE transactions on pattern analysis and machine intelligence*, vol. 33, no. 12, pp. 2341–2353, 2011.

[43] J. M. Morel, A. B. Petro, and C. Sbert, "A PDE formalization of Retinex theory," *IEEE Transactions on Image Processing*, vol. 19, no. 11, pp. 2825–2837, 2010.

Michela Lecca received her Master Degree in Mathematics from the University of Trento, Italy. She is currently researcher at the Research Unit Technologies of Vision of Fondazione Bruno Kessler - Center for Information and Communication Technology, in Trento, Italy. Her research interests include color image processing, color constancy and correction, automatic object recognition, semantic image labeling, and hardware oriented image analysis. Michela took part to the program committee of many editions of Applied Computing Machine Symposia, and carries on a reviewing activity for several international conferences and journals. She is a member of Italian Association for Computer Vision, Pattern Recognition and Machine Learning (IAPR-CVPL) and of the Gruppo Italiano del Colore - Associazione Italiana Colore (GdC-AIC).