

Statistical Aspects of Space Sampling in Retinex models

Gabriele Gianini; Università degli Studi di Milano, Computer Science Department, via Bramante 65, 26013 - Crema (CR), Italy

Abstract

A distinguishing element of the different models in the Retinex family, is the process by which the space surrounding the target pixel is explored. Several models use a process define in terms of repeated sampling, prescribe to compute a synthesis quantity out of each pixel sample and finally to take an average of those quantities to determine the output. For instance, in MI-Retinex each sample is define by a memoryless random walk, connecting a randomly chosen reference pixel to the target, while in RSR the sample is define by a random set of points centered on the target; in both the averaged quantity is a suitably computed sample maximum. Here we discuss the advantages that arise from a probabilistic reformulation of each of those statistical sampling process, with reference to two recently formulated Retinex models: ReMark and QBRIX. In those two algorithms the sampling-computing-and-averaging process is replaced by the direct calculation of the sampling means of the synthesis quantity, out of the whole population of possible samples. In QBRIX – inspired to RSR – this corresponds to computing, for each pixel, the probability that it becomes the maximum of a spray: overall it reduces to the determination of a high quantile of the histogram of pixel intensities. In ReMark – based on MI-Retinex – this corresponds to computing, for each pixel, the probability that it represents the last point of reset, before the random walk meets the target: this entails rethinking the process as a Markov Process and computing its absorption probabilities, by solving a linear system. In this work, we compare informally the two approaches and argue that reasoning in terms of population models can bring new insight into the features of the distinct Retinex variants and highlight connections and differences among mathematical models.

Introduction

The Retinex model by Land and McCann [1, 2] is a model of the Human Vision System (HVS) consisting in some well established core elements and in some other elements, which, along the years, have undergone several different formalizations. Among the variants of the model, some are based on a repeated sampling and averaging process: here we focus on two specific sampling-based models, MI-Retinex [3] and RSR [4].

Repeated sampling is a process that can either be enacted, as in MI-Retinex and RSR, to compute the average of a sample define quantity f (e.g. the maximum intensity of the sample), or be modeled probabilistically. In the second case one can, in principle, obtain the limiting value of the average by computing the mean of f over the population of all the possible samples. Hereafter we refer to those algorithms that prescribe to enact the repeated sampling as *sampling based models*, and to those algorithm which attempt to compute the mean of the sample based quantity out of the population of all the possible samples as *population based models*, the short for sampling population models, or sampling distribution models.

In this note we describe MI-Retinex and RSR and then review their probabilistic formalizations, respectively ReMark [5] and QBRIX [6]. The present work complements the formalizations in [5] and [6] by developing the discussion at an informal level and by leveraging intuitive arguments. We focus on some of the relationships among models and not on their ability to reproduce the HVS behavior.

We are aware that, in the case of ReMark, interesting links with existing models could arise from the continuum-space limit of the corresponding Markov Chains, and from the implied PDE models; however, highlighting the differences with the existing PDE models would require a careful and long discussion: due to the complexity of the subject and to space limitations, this topic is left outside the scope of the present paper.

The paper is structured as follows: first we recall the defining elements of the Retinex model; then we recall the definition of MI-Retinex and of its "point spray" based approximation, RSR; subsequently we present, in order of increasing complexity, the population based representations of RSR (QBRIX), and the population based representation of MI-Retinex (ReMark). In each case we point to merits and draw-backs of the sampling- and of the population-based representations.

Retinex models

Among the core components of the Retinex algorithm formalizations are the chromatic channel separation, the reset mechanism and the thresholding mechanism. The first component is motivated by the observation that the HVS operates in three independent retinal-cortical systems, processing respectively the low, middle and high frequencies of the visible spectrum: algorithmically it translates into the prescription that each channel has to be processed independently of the others. The reset and thresholding mechanisms are motivated by the observation that color sensation, and consequently color appearance, are not based on the color stimulus at the point, but rather on the relative spatial arrangement of the stimuli in the observed scene [7]: both intensities and gradients take part into the determination of the output. Each Retinex system attenuates smooth changes in intensity – arising for instance from gradients of the illuminant – and enhances sharp changes, i.e. the edges; furthermore, each Retinex system estimates the appearance of a point by relating the value in the point to a *local reference white* found within the image and rescales all the intensities of the region with respect to that reference. Algorithmically, in the original path based formulation of Retinex, these observations were rendered by a *path-level* mechanism, called *reset*, which would change the current reference white whenever the path exploring the image met a new maximum, and by a *step-level* mechanism, called *thresholding*, that would discount small intensity gradients found along the path.

Beside those well established elements, the early Retinex formulations contain components which have been often reinter-

preted and given different algorithmic formalizations: the most relevant of those components is the space sampling process, i.e. the process by which the space surrounding the target pixel is explored. This process contributes in a decisive way to the determination of a desirable feature of the HVS models, the so called *locality*, which consists in the fact that the correction at a point depends strongly on the input intensities of the closer regions and weakly on those of farther regions. We return on locality later.

The space sampling process proposed in the early Retinex formulations consisted in a pixel (or region) sampling guided by a path whose deterministic or random character depended on the implementations. Some relevant variants defined by means of path-based sampling are MI-Retinex [3], Termite-Retinex (TR) [8], Energy-driven Termite Retinex (ETR) [9] and the algorithm by Montagna and Finlayson [10].

In MI-Retinex [3] the paths are defined as memoryless random walks starting from a randomly chosen pixel and stopping the first time the path meets the target pixel. In Termite Retinex [8] the paths are defined by stigmergic random walks: each path proceeds randomly but chooses preferably the regions less visited by the previous paths (the history of the previous paths is captured by a "pheromone landscape" to which every path equally contributes its "pheromone trail"). In ETR [9] the paths are not determined by a local search, but by a search in the path space: the algorithm selects, from the collection of all paths connecting a random starting point to the target, those that fulfil a set of specific constraints: they are local minima of a complex energy functional, taking into account the pheromone landscape, the distance from the target and the intensity gradients. Finally the model by Montagna and Finlayson [10] uses a pseudo-Brownian path, constrained so as to guarantee a lower bound to the number of visits to each pixel, and an even average number of visits per pixel.

Besides the formalizations of the paths, other defining elements contribute to the diversity of Retinex variants: the combination of mechanisms is also open to different choices. For instance in MI-Retinex the quantity computed along the path is used to update the target pixel only; on the contrary, in the more classical version by Land and McCann [1, 2] and in the version by Montagna and Finlayson, each pixel along the path is updated with the information collected by the path up to that point.

MI-Retinex

In MI-Retinex (as in TR and ETR) the image exploration process is based on repeated statistical sampling: each random walk connecting the random starting point to the target provides a sample of pixels; MI-Retinex prescribes to compute a synthesis quantity out of each path and eventually to take an average of those synthesis quantities to determine the output value at the target pixel. The synthesis quantity from each sample is computed (in each chromatic channel) by applying the reset and the thresholding mechanism to the products of the ratios of intensities of subsequent pixels along the path: the thresholding consists in forcing the ratio to one when it is within some small threshold from the unit; the reset consists in forcing the product of the ratios to one whenever it exceeds one (for a formal account see [5]).

Sometimes – since it has been observed [3] that the suppression of gradients has a minor impact on the output from most natural images – the MI-Retinex model is implemented without thresholding (we called this variant Reset-only MI-Retinex). In

that case the computation of the synthesis quantity out of a path can be shown to reduce to the computation of the maximum intensity of the set of pixels visited by the path. This quantity is averaged over all the paths and used as a new reference white level for the target, with respect to which the input is rescaled.

It is worth pointing out that applying the reset mechanism is equivalent to enforcing a memory-loss to the path every time it finds a new maximum: in the Reset-only version, the value contributed to the target for the computation of the reference white level equals the intensity of the last reset point.

Locality

The combined effect of path-based exploration and memory loss at reset points determines a desirable feature: the "locality" of filtering. A filter is local if the information received and used at the target point comes preferably from nearby points. If the information from a point is used equally over the whole image, the correction is global. A prototypical example of global algorithm is the von Kries filter, a.k.a. Scale-by-Max algorithm, which looks for the maximum intensity present in the image and uses that single value as the reference white level for all the pixels.

MI-Retinex is local by construction. In MI-Retinex locality is due to two defining factors: a geometric factor and a mechanism related factor. The geometric factor depends on the fact that the algorithm uses those random paths that have an end at the target and another at a uniformly randomly chosen point of the image: as a consequence the points visited by the paths will be located more frequently in proximity of the target (this locality determining mechanism is image-independent).

The second locality determining factor is built-in in the reset mechanism: due to the memory loss enforced by the reset mechanism, the information arriving to the target is determined by the closest reset point met by the path (this locality determining mechanism is image-dependent).

RSR

The Random Spray Retinex (RSR) [4] algorithm builds on the modeling features formalized by the Reset-only MI-Retinex and achieves a higher efficiency, by trying to retain most of the qualitatively desirable features of the path-based algorithm.

The key observation of RSR is that one could replace the expensive process of random walk generation by a more efficient one. The MI-Retinex process is inefficient in that the local search of the target can take a very long time and in that intermediate pixels visited several times do not bring extra information. The authors of [3] propose to substitute each random path by a set of points – called spray – generated so as to loosely mimic the path sampling: they choose to define the spray as a set of points sampled from a radially symmetric *sampling profile* (formally, the process corresponds to an isotropic non-homogenous Poisson point-sampling process with intensity decreasing as the distance from the target increases).

In RSR, locality is obtained and controlled by the geometry of the sampling profile i.e. by the speed of decrease as a function of the distance from the target (for a recent variant of RSR see also [11]).

From sampling-based to population-based representations

When the values obtained by computing a sample based function f repeatedly over different independent samples are aggregated by taking their arithmetic average ($\langle f \rangle$), one can think of the resulting quantity as of an estimate of the true mean μ_f of the sample based function. Often the latter can be computed with precision directly out of the reference population, without resorting to sampling-computing-and-averaging process (which can yield only approximate results, a potential source of noise).

Consider the following illustrative example. Take a random uniform variable with values $i \in [0, 1]$ (a density over the continuous standard unitary interval). Consider a fixed sample size n . Choose n values randomly and independently, so as to obtain the sample $\{i_k^{(s)}\}_{k=1}^n$ (the exponent s indexes the sample), finally compute the maximum of the sampled values:

$$m_s \equiv \max \left\{ i_k^{(s)} \right\}_{k=1}^n .$$

Pick a large but finite number N of fixed size samples and repeat the above procedure, i.e. compute, for each sample, the maximum value m_s ; finally average the sample maxima:

$$\hat{\mu}_m = \frac{1}{N} \sum_s m_s .$$

This process can be thought as the computation of the estimate of the true mean of the sampling maximum (a quantity which could be in principle recovered only in the limit of an infinite number N of samples). Note, in passing, that running this estimate procedure several times (each time with N samples) would yield each time a different value: the estimates fluctuate around the true value μ_m .

On the other hand, this process can be modeled probabilistically, considering the population of *all* the possible samples. The probability (density) that a given value of i is the maximum of a sample of n randomly and independently chosen values from a uniform density is $g_{max}(i) = n i^{(n-1)}$, and its center of mass is

$$\mu_m = \frac{n}{n+1} .$$

This corresponds to the exact value of the mean of the sampling maximum, recovered by direct computation out of the population.

This toy example illustrates the relation between a sampling-based algorithm to estimate the sampling maximum and the corresponding population-based computation.

Sampling-based and population-based methods are complementary to one another and are endowed with advantages and drawbacks. Sampling-based methods are easy to formulate, offer a direct perception of possible mechanism variants and mechanisms approximations (consider for instance the relation between MI-Retinex and its RSR approximation), however they are intrinsically noisy. MI-Retinex and RSR can be considered sampling-based algorithms having the purpose of computing an estimate of a given quantity (the white-reference level for a target pixel). As such they can be affected by noise: repeating the procedure for nearby pixels of similar intensity can yield perceptibly different corrections, i.e. introduce noise into the filtered image. The price to pay to avoid this effect amounts to using a large number N of samples (the precision of the estimate scales with the usual square

root reciprocal law). So, both MI-Retinex (as the other path based algorithms) and RSR are need in general a non-negligible computational effort to make the noise non perceivable.

Population-based methods, on the other hand are intrinsically noise-free, however they are in general more difficult to formalize and often entail a complex implementation. An advantage of population-based models is that their form is, in a sense, closer to the description of the output, so that they can give useful hints about the outcome and suggest useful approximations. Moreover, when the population-based model is an exact translation of the sampling-based model, the parameters of the latter are in direct correspondence to the parameters of the former, so one can more directly appreciate the effect of a parameter value change. Finally such models can point to other high level models (e.g. links among analytical models) and solution/computation tools (or exclude that some tools can be used in the modeling).

QBRIX: a population model inspired to RSR

The general qualitative characters of the population-based representation of RSR are rather straightforward to obtain from the sampling-based model. RSR depends on a small set of parameters: the shape of the sampling profile the number of points n in a sample (a spray), and the number N of sprays used for each target pixel. Normally, with RSR, one filter an image by tuning the parameters until a satisfactory output is obtained.

In a population model the parameter N can be thought as set to $+\infty$, so in a population model of RSR one is left only with two parameters: n and the *sampling profile*. In order to appreciate the qualitative features of the population-based representation of RSR let us consider a *flat* sampling profile (the spray pixels are sampled with the same probability from any point of the image), and consider n as the only variable parameter. We call this model global-RSR, or gRSR (due to the flat sampling profile the influence of the pixels over the target does not depend on their distance from the target): gRSR is a 1-parameter model.

Now let us reconsider the above toy example, i.e. imagine taking samples of size n out of a uniform density: this density can represent the idealized intensity histogram of a toy image whose intensities are uniformly distributed over the interval $[0, 1]$ (notice that, since the spatial sampling profile is flat the geometric arrangement of the intensities in this image is irrelevant). However, let us introduce a key difference: out of each sample let us compute, not the maximum intensity, but its reciprocal $1/m_s$ and finally take the average

$$\hat{\mu}_{\frac{1}{m}} = \frac{1}{N} \sum_s \frac{1}{m_s} .$$

This formulation corresponds to running $gRSR(n)$ over the input image: indeed the algorithm define the output value \hat{o}_{gRSR} of a pixel intensity based on the input value i of the pixel intensity as $\hat{o}_{gRSR} = i \cdot \hat{\mu}_{\frac{1}{m}}$, or, equivalently, using the reference white level \hat{w} , define by the harmonic mean $\frac{1}{\hat{w}} = \hat{\mu}_{\frac{1}{m}}$, as

$$\hat{o}_{gRSR} = \frac{i}{\hat{w}}$$

Now let us think of $\hat{\mu}_{\frac{1}{m}}$ as to an estimate of the "true" mean $\mu_{\frac{1}{m}}$ of the sample-based function $1/m$: thanks to the fact that we are

using the above define toy image, it very easy to work out the "true" mean and the corresponding "true" reference white level $w = 1/\mu_{\frac{1}{n}}$. It turns out that (for the details of the computations see [6]) the gRSR reference white level over this image is

$$w = 1 - \frac{1}{n}.$$

Based on this value, the output o_{gRSR} of any target pixel of the toy image can be computed out of the target input intensity i as

$$o_{gRSR} = \frac{i}{w} = i \left(1 - \frac{1}{n}\right)^{-1}.$$

An important observation at this point is that w corresponds to a *high quantile* of the uniform distribution (it is the complementary quantile $CQ(\frac{1}{n}) = F^{(-1)}(1 - \frac{1}{n})$ of the fraction $\frac{1}{n}$) and here its value depends directly upon the only parameter of the gRSR model, i.e. n . Specifically, the higher is n , the closer w is to 1 (the global maximum of the image). This is the starting point of the QBRIX algorithm. Let us call $gQBRIX$ the global version and $lQBRIX$ the local version of the algorithm: the equations above can be rewritten as

$$o_{gQBRIX} = \frac{i}{w} = \frac{i}{CQ(\frac{1}{n})} = i/CQ(q),$$

where q can be used as a parameter of $gQBRIX$, in place of $1/n$.

Of course in general the determination of w is not straightforward when the input image is not as simple as this one. Nonetheless, given the image intensity histogram, one can tune the filtering by changing q until the output image is satisfactory. Furthermore, when the sampling profile is non-flat one can define a local version of the algorithm, $lQBRIX$, by which the program receives in input, for each target pixel, not the original image intensity histogram, but rather an histogram obtained by weighting the intensities based on the distance from the target and according to the spatial sampling profile. We stress that $lQBRIX$ is not the exact population-based model corresponding to RSR, but rather an approximation obtained on the basis of first principles: whatever the value of n in RSR, the reference white level will turn out to be a high quantile of the distance-weighted intensity distribution.

More details and examples are given in [6], where the quantile based algorithm is defined and the effect of the approximation is discussed more thoroughly, also considering the fact that the RSR algorithm prescribes to include the target pixel in each spray.

Let us conclude the discussion about RSR by pointing out that, thanks to the quantile view of RSR/QBRIX, one can establish a relation between known models: the von Kries/Scale-by-Max filter, for instance, can be seen as a special case of $gQBRIX$, using the highest quantile.

In synthesis, the QBRIX algorithm is a population-based model providing an approximation of the exact population-based mapping out of a sampling-based model (in other cases an exact mapping can be used, as illustrated below for MI-Retinex): the insight provided by the new model suggests to use a new parameter, closer to the description of the output, a high quantile value of the intensity histogram (here indicated by $1 - q$): this implies that it is the fraction q of top intensity pixels (weighted by their distance from the target), that determines the reference white level.

ReMark: the population model of MI-Retinex

MI-Retinex samples the image by repeatedly generating memoryless random paths, connecting a randomly chosen point of the image to the target: each path brings to the target a contribution computed along the path by means of the threshold-and-reset functional. For the sake of simplicity we limit the discussion to the Reset-only version of MI-Retinex, whose distinguishing element consist in the fact that the quantity computed along the path reduces to the maximum pixel intensity.

The Reset-only MI-Retinex prescribes to consider a randomly chosen starting point in the image, called (initial) reference, to build a memoryless random walk starting from that point and to stop the walk the first time it meets the target: this process determines a sample, consisting in the collection of pixels visited by the walk (each pixel can be visited several times, except the target, but to the outcome of the computation the number of visits has no relevance). We indicate by n the number of distinct visited pixels, and the collection of their intensities by $\{i(x_k^{(\gamma)})\}_{k=1}^n$, where γ is the index referring to the individual walk. Then (as in RSR, which was inspired to this algorithm) the Reset-only MI-Retinex prescribes to compute the maximum of the sample of values;

$$m_\gamma \equiv \max \left\{ i \left(x_k^{(\gamma)} \right) \right\}_{k=1}^n.$$

The path based pixel sampling and the following maximum computation are repeated a large number N of times; finally the reciprocals of the sample maxima are averaged to obtain the reference white value \hat{w} from

$$\frac{1}{\hat{w}} = \frac{1}{N} \sum_{\gamma} \frac{1}{m_\gamma}$$

and the output, from the target intensity i , as $o_{MIRetinex} = i/\hat{w}$.

To provide a population-based translation of this process, one has to compute a weighted sum, and weight the intensity of each pixel by the probability that it becomes the maximum of a reference-to-target path, then take the result as the white reference level to compute the output. The computation of the relevant probability is non-trivial. In [5] this is accomplished by modeling the joint process – consisting in memoryless random walk over the image and in the the path functional computation – as a suitable Markov Process. Those models allow to consider at the same time two levels of description: in our case, the description of the walk displacement in the 2D geometric space of the image, and the evolution of the information carried by the walk (the maximum intensity found up to the current step).

To represent the latter, in [5] a 3D embedding is used: two dimensions are reserved to the displacement over the image (represented as a diffusive process) and the third to the displacement in the 1D space of the carried information as illustrated in Fig. 1. Notice that to each reset point of the walk (a point where the path has found a new maximum of intensity), corresponds a level jump in the 3D representation: those points are seen as sinks of the Absorbing Markov Chain by a level and as sources by another.

Notice also that the target is represented by different points: the distinct *target representatives* are situated along the same vertical: if the path meets one of those points it stops and contributes to the target the value of the intensity level on which it is walking. By using Absorbing Markov Chains one can compute for every source point, situated on its 2D coordinate and at its level of

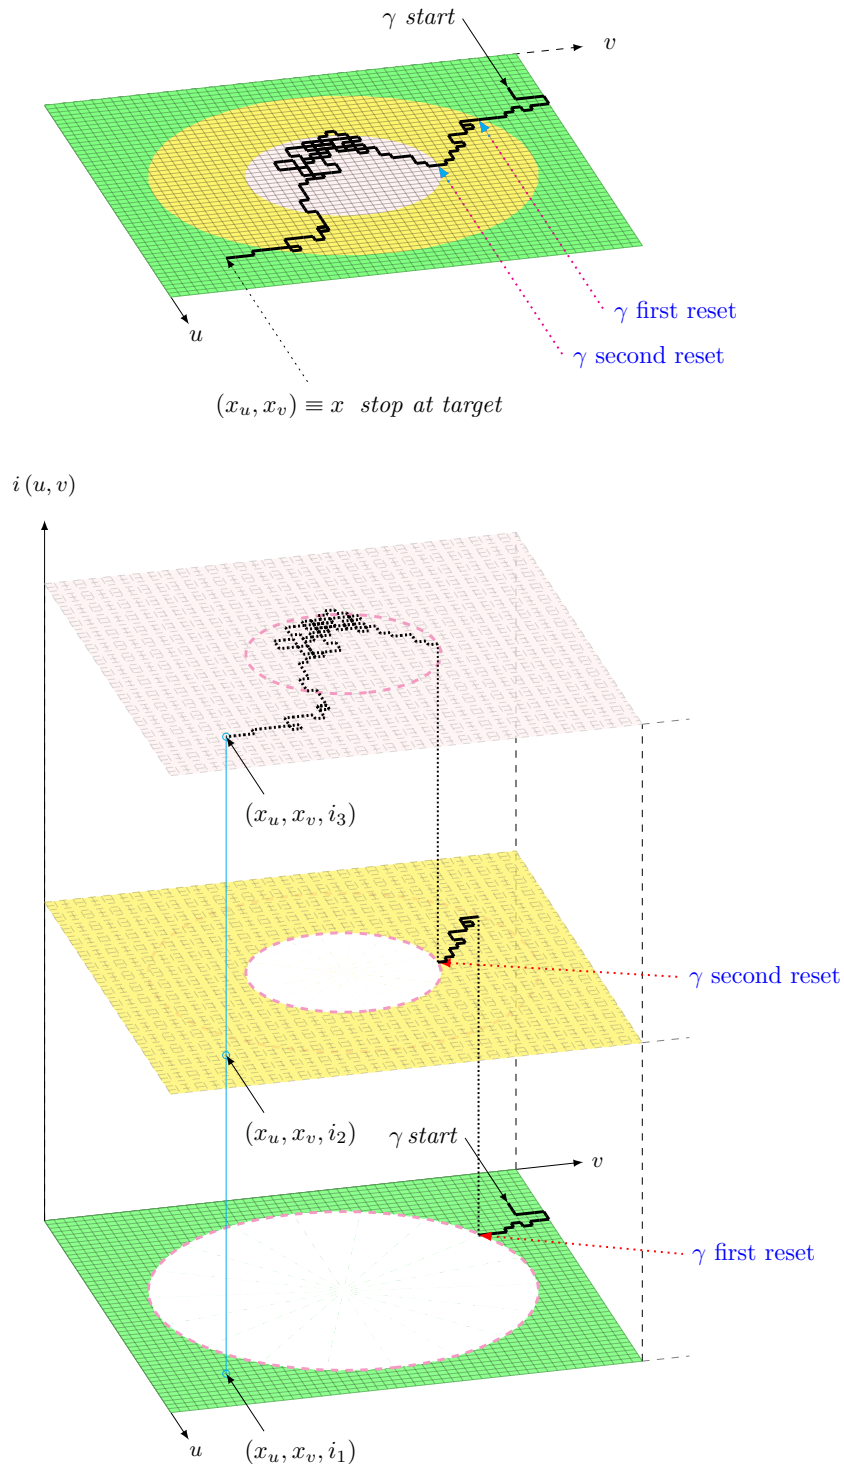


Figure 1. An example random walks from the 2D space representation (top) to the 3D Markov Chain representation (bottom). The different colors in 2D space represent different intensities of a single chromatic channel: the closer the center of the image, the higher the intensity of the region. In the Reset-only MI-Retinex [3] the 2D random walk carries the information about the latest intensity maximum found in exploring the image. Every time a new maximum is found, the carried information is updated: the point of the image where the update takes place is called reset point. This process can be modeled as a 3D Markov Process: two dimensions are used to model the spatial diffusion, the third dimension to model the carried value: reset points correspond to upward jumps in the 3D representation. The target pixel x is pictured by a vertical array of target representatives: each representative absorbs and stops the walk. The fraction of walkers absorbed by target representative located at a level i corresponds to the probability p_i that that level of intensity contributes to the target. ReMark [5] uses those probabilities as weights to compute $1/w \equiv \sum_i p_i/i$ and takes w as new reference white level for a target x ; then it computes the output $o(x)$ from the input $i(x)$ as $o(x) = i(x)/w$.

intensity what is the probability that a random walk starting from there will be absorbed at in correspondence of a representative of the target. One can also aggregate the sources from the same starting level: the Absorbing Markov Chains provide the fraction of the population of "walkers" absorbed at each representative point and use that fraction as a weight. Finally, by weighting the reciprocal of the absorption levels by that weight one obtains the desired reference white value for the target.

The actual computation consists in the inversion of a large matrix, the so called Fundamental Matrix of the Absorbing Markov Chain, derived from the transition matrix. Since the matrix is sparse (each state is connected only to few neighbors) one can manage to process even large images, despite the large number of states involved (bound superiorly by the number of pixels times the number of levels). Still the computation of this exact mapping of MI-Retindex is rather expensive. Nonetheless, if one aims at producing noise-free filtered images, the algorithm, called ReMark, turns out to be more efficient than the strictly memory-less random walk of MI-Retindex [3] (further details in [5]).

As usual, the population-based model is characterized by a lower number of parameters w.r.t. the sampling-based model. Indeed, ReMark is defined mainly by the spatial distribution of the sampling profile used for choosing the starting point of the walk. Instead, the corresponding sampling-based model, MI-Retindex, is controlled by the following additional parameters: the number N_{ref} of reference pixels sampled for each individual target; the number N_{walks} of walks for each reference-target pair; the maximum length S of a walk, i.e. the maximum number of reference-to-target steps that a random walk is allowed to take. The first and the second parameter control the noisiness of the output, the third is used to control the execution time: in practical implementations it is necessary to set this cutoff threshold because the convergence of random walk local search for the target is exceedingly long. Ideally in the corresponding population model, i.e. in ReMark, all the three parameters are set to $+\infty$.

The comparison of this ideal case can be used to probe the limits of the necessarily finite values of some sampling model parameters. For instance, an observation that emerged in [5] points to the fact that the parameter S , apparently a nuisance parameter, has a perceptible role in controlling locality: comparing the ReMark output to the MI-Retindex output for different values of S one can see that it takes a very high cutoff value S in the number of steps (of the order of 10^5 for a 100×100 image) to approximate adequately the ideal output.

Conclusions

Sampling-based Retindex models and population-based Retindex models are complementary to one another – metaphorically, as the individual tree view vs. forest view – and are each endowed with merits and drawbacks. Original insights into the modeled phenomenon can be provided by each type of model individually and by their mutual comparison.

Acknowledgements

This work was partly supported by the Gastwissenschaftlerprogramm of the University of Passau and by the INSA of Lyon through the program CMIRA2014 (Subvention n.14.007051). The author would like to express his sincere appreciation for the fruitful discussions with Alessandro Rizzi and Michela Lecca.

References

- [1] E. H. Land and J. J. McCann, "Lightness and Retindex Theory," *Journal of the Optical Society of America*, vol. 61, pp. 1–11, Jan. 1971.
- [2] E. H. Land *et al.*, *The retindex theory of color vision*. Scientific American., 1977.
- [3] E. Provenzi, L. De Carli, A. Rizzi, and D. Marini, "Mathematical definition and analysis of the retindex algorithm," *JOSA A*, vol. 22, no. 12, pp. 2613–2621, 2005.
- [4] E. Provenzi, M. Fierro, A. Rizzi, L. De Carli, D. Gadia, and D. Marini, "Random spray retindex: a new retindex implementation to investigate the local properties of the model," *Image Processing, IEEE Transactions on*, vol. 16, no. 1, pp. 162–171, 2007.
- [5] G. Gianini, A. Rizzi, and E. Damiani, "A retindex model based on absorbing markov chains," *Inf. Sci.*, vol. 327, pp. 149–174, Jan. 2016.
- [6] G. Gianini, A. Manenti, and A. Rizzi, "Qbrix: a quantile-based approach to retindex," *JOSA A*, vol. 31, no. 12, pp. 2663–2673, 2014.
- [7] J. Albers, *Interaction of Color*. New Haven: Yale University Press, 1975.
- [8] G. Simone, G. Audino, I. Farup, F. Albregtsen, and A. Rizzi, "Termite retindex: a new implementation based on a colony of intelligent agents," *Journal of Electronic Imaging*, vol. 23, no. 1, pp. 013006–013006, 2014.
- [9] M. Lecca, A. Rizzi, and G. Gianini, "Energy-driven path search for termite retindex," *J. Opt. Soc. Am. A (JOSA A)*, 2016.
- [10] R. Montagna and G. D. Finlayson, "Constrained pseudo-brownian motion and its application to image enhancement," *JOSA A*, vol. 28, no. 8, pp. 1677–1688, 2011.
- [11] M. Lecca and A. Rizzi, "Tuning the locality of filtering with a spatially weighted implementation of random spray retindex," *J. Opt. Soc. Am. A*, vol. 32, pp. 1876–1887, Oct 2015.

Author Biography

Gabriele Gianini has received his M.Sc. Degree in Physics (1992) and his Ph.D. in Physics (1996) from the Università degli Studi di Pavia (Italy). Since 2005 is Assistant Professor at the Department of Computer Science of the Università degli Studi di Milano (Italy), where he is lecturer of Probability and Statistics. He has held visiting positions at a number of international institutions, including INSA de Lyon (France), University of Passau (Germany), CERN (Geneva, Switzerland), Fermilab (Chicago, US), CBPF (Brasil). He has been Adjoint professor at the Free University of Bolzano (Italy) from 2005 to 2012. Among his research interests: image processing, quantitative modeling of processes, statistical and soft computing techniques. From 1992 to present has co-authored over 150 papers published in internationally refereed journals and international conference proceedings.