https://doi.org/10.2352/ISSN.2169-2629.2018.26.86

An alternative multiscale framework for variational perceptually-inspired contrast enhancement of color images

Baptiste Mazin (independent researcher) and Edoardo Provenzi (Institut of Mathematics of the University of Bordeaux, France)
baptiste.mazin@gmail.com, edoardo.provenzi@u-bordeaux.fr

Abstract

We present a general multiscale strategy for perceptually-inspired contrast enhancement of color images. The idea behind this methodology comes from a recent wavelet-based variational framework for contrast intensification. We will show that the equations for the wavelet coefficients coming from the variational setting can be re-written in a more general multi-resolution framework, where the only requirement is the existence of an approximation and a detail layer at each different scale. In particular, we will show that a Laplacian pyramid implementation of the variational algorithm performs faster and better than the wavelet-based one. These results open the possibility to efficiently apply the contrast enhancement equations also to video sequences.

Introduction

In the papers [3, 6, 2] the authors introduced a variational framework where several models of perceptually-inspired color enhancement, e.g. [13, 10, 11, 12], could be embedded and regarded as special instances of a more general theory. A detailed algorithmic description of these variational models can be found in the paper [4] and a theoretical overview can be read in the book [8].

In [7, 9], the variational framework has been recast in the setting of wavelets. The equations corresponding to this new formulations happened to be much easier to implement than those referring to the original variational framework. However, the constraints that have to be respected in order to guarantee the stability and convergence of the wavelet-based algorithm, and also to avoid the introduction of the typical wavelet-like artifacts, are quite restrictives and, in certain cases, they can noticeably reduce the enhancement power of the method.

With the aim of reducing these drawbacks, we have studied the possibility to re-formulate the wavelet-based variational framework by keeping intact its intrinsic multi-resolution nature and the simple Euler-Lagrange equations, while using approximation and coefficient details coming from a multiscale pyramid alternative to wavelets. We will show that this reformulation is actually possible and it leads to a better performing algorithm, as we prove with a Laplacian pyramid implementation.

The plan of this work is the following. In section we will introduce the concept of perceptual functionals for contrast enhancement of color images and their wavelet counterparts. In section we will extend the wavelet-based variational framework to a more general multiscale one, discussing, in particular, a Laplacian pyramid implementation. Tests and comparisons will be discussed in section and perspectives for further applications will be given in the final section .

Perceptually-inspired contrast enhancement in the spatial and wavelet domain

In the papers [6], the authors shown that the only class of energy functionals complying with all four basic phenomenological characteristics of the Human Visual System (HVS from now on) is the following:

$$E_{\mu,I_0,w,\varphi}(I) = D_{\mu,I_0}(I) + C_{w,\varphi}(I), \tag{1}$$

where

$$D_{\mu,I_0}(I) = \int_\Omega \left[\alpha\left(\mu \log\frac{\mu}{I(x)} - (\mu - I(x))\right) + \beta\left(I_0 \log\frac{I_0}{I(x)} - (I_0 - I(x))\right)\right] dx, \tag{2}$$

and

$$C_{w,\varphi}(I) = \iint_{\Omega^2} w(x,y)\,\varphi\left(\frac{\min\{I(x),I(y)\}}{\max\{I(x),I(y)\}}\right) dxdy \tag{3}$$

where:

- $I : \Omega \to [0,1]$ is any chromatic channel R,G,B of the objective image function;
- I_0 is any given chromatic channel of the original image function;
- $C_{w,\varphi}(I)$ is the contrast enhancement term, its minimization leads to a local non-linear intensification of spatial contrast;
- $D_{\mu,I_0}(I)$ is the entropic dispersion term, its minimization prevents the objective image to depart too much from both the average and the original image intensities;
- μ is the average value of I_0;
- $\alpha, \beta > 0$ are real coefficients which control the attachment to μ and to the original image function values I_0, respectively;
- $w : \Omega \times \Omega \to \mathbb{R}^+$ is a spatial kernel, which depends only the Euclidean distance $\|x - y\|$ among two generic pixels $x, y \in \Omega$ and it is monotonically decreasing with the distance itself, i.e. $w(x,y) := g(\| x - y \|)$ where $g : \mathbb{R}^+ \to \mathbb{R}$ is monotonically decreasing;
- $\varphi : [0,1] \to \mathbb{R}$ is a monotonically increasing differentiable function.

To understand the motivation for the analytical shape of the contrast term $C_{w,\varphi}(I)$ let us consider the basic image formation model, i.e. $I(x) = \rho(x) \cdot \lambda$, where $\rho(x)$ represents the reflectance of a point x and λ represents the illuminant (supposed to be constant all over the scene). Since $C_{w,\varphi}(I)$ is written in terms of a ratio, it is evident that it is independent with respect to illuminant changes, thus it is coherent with the color constancy feature.

Moreover, as proven in [6], this analytical form is the only one in which color constancy and Weber's law can be combined.

The function φ represents a degree of freedom and it is chosen to be monotonically increasing not to reverse contrast between pixels.

Finally, notice that the minimization of $C_{w,\varphi}(I)$ indeed induces a contrast enhancement. In fact, the ratio $\min\{I(x),I(y)\}/\max\{I(x),I(y)\}$ is minimized when the lowest value between two pixels is decreased and the highest is increased, which, of course, corresponds to an intensification of contrast. The enhancement is spatially local due to the presence of the weighting function w. Typically w is a Gaussian kernel with center in x, its standard deviation σ can be set by a user to increase or decrease the locality of contrast enhancement. Small values of σ push the effect towards *sharpening*, large values of σ instead push towards a global enhancement.

Let us now discuss the dispersion term. Its choice has been guided by dimensional coherence with the contrast term, which has dimension 0 with respect to I. The easiest meaningful candidate is the *entropy functional* D_{μ,I_0}, whose minimization produces a reduction of entropy, i.e. disorder, around the average value μ (which can be different in each chromatic channel) and around the original image I_0. This last attachment is introduced to avoid an excessive departure from the original intensity values and can be modulated via the coefficients ratio α/β.

In [9], the authors re-wrote the functional (1) in the framework of wavelet theory. The main idea behind their work is to represent local contrast via a functional of detail wavelet coefficients, balanced by adjustment terms as in spatial framework just discussed. Let us first consider the effect of adaptation to the average level. If we were dealing with Fourier transforms, the average image intensity value μ would be proportional to the zero-th Fourier coefficient, but in the wavelet domain this direct correspondence is not available. However, since the coefficients $\{a_{j,k}, k \in \Omega\}$ represent the image approximation at the scale 2^j, a natural analogue of μ in the wavelet framework at the scale 2^j is represented by $\overline{a}_j \equiv \frac{1}{|\Omega|}\sum_{k\in\Omega} a_{j,k}$, i.e. the average approximation coefficient. $\overline{a}_j$ can be considered per-channel, or averaged over the three chromatic channels.

Let us also underline that we only need to modify the approximation coefficients of the coarsest scale, since this modification will be propagated to finer scales via the sum of the modified detail coefficients, as we will formalize at the end of this section. We implement the adaptation to the average value at the coarser scale balanced by the attachment to the original values through this convex linear combination

$$a_{J,k} \equiv \alpha\overline{a}_J + (1-\alpha)a^0_{J,k}, \tag{4}$$

where $\{a^0_{J,k}, k \in \Omega\}$ is the original sequence of approximation coefficients at the scale 2^J and $\alpha \in [0,1]$ is a suitable weight coefficient, whose influence will be discussed in test section. The bigger α, the strongest the attachment to the average value $\overline{a}_J$, and viceversa.

Let us now pass to the wavelet-based version of the perceptual contrast functional, which is the following:

$$\mathscr{C}_{w_j,\varphi,\{a_{j,k}\}}(\{d_{j,k}\}) = \sum_{k\in\Omega} w_j\,\varphi\left(\frac{a_{j,k}}{d_{j,k}}\right), \qquad 2^J \geq 2^j \geq 2^{L+1}, \tag{5}$$

where $\{a_{j,k}, k \in \Omega\}$ and $\{d^\ell_{j,k}, k \in \Omega\}$, $\ell = H,V,D$ are the approximation and (horizontal, vertical and diagonal) detail wavelet coefficients of the dyadic level 2^j, 2^J and 2^L are the coarsest and the finest level, respectively (see, e.g. [5] for more information about the wavelet framework). To simplify the notation, we have eliminated the superscript ℓ in the detail coefficients, by making the implicit assumption that the operations are repeated on the three kinds of detail coefficients. With the notation $k \in \Omega$, we will implicitly assume a column-wise ordering of Ω, the spatial support of the image, so that Ω can be seen as a finite subset of $\mathbb{Z}$. w_j are positive coefficients that permit to differentiate the contrast enhancement action depending on the scale 2^j and $\varphi : [0,\infty) \to [0,\infty)$ is a differentiable monotonically increasing function such that $\varphi(r) \to +\infty$ as $r \to \infty$.

We underline that, since changing the sign of a wavelet coefficient can result in drastic modifications of an image, we will modify only the *absolute value* of the wavelet coefficients, restoring the original sign at the end of the procedure.

In order to prevent an excessive magnification of the original detail coefficients, whose absolute value is denoted with $d^0_{j,k}$, a conservative term should be introduced. Again, to maintain dimensional coherence with $\mathscr{C}_{w_j,\varphi,\{a_{j,k}\}}$, an entropic dispersion functional is a suitable choice:

$$\mathscr{D}_{d^0_{j,k}}(\{d_{j,k}\}) \equiv \sum_{k\in\Omega}\left[d^0_{j,k}\log\frac{d^0_{j,k}}{d_{j,k}} - \left(d^0_{j,k} - d_{j,k}\right)\right], \tag{6}$$

$2^J \geq 2^j \geq 2^{L+1}$.

Combining these two effects one can define the energy functional that realizes local contrast enhancement as $\mathscr{E}_{w_j,\varphi,\{a_{j,k}\},d^0_{j,k}} = \mathscr{C}_{w_j,\varphi,\{a_{j,k}\}} + \mathscr{D}_{d^0_{j,k}}$, i.e.

$$\mathscr{E}_{w_j,\varphi,\{a_{j,k}\},d^0_{j,k}} \equiv \sum_{k\in\Omega}\left[w_j\varphi\left(\frac{a_{j,k}}{d_{j,k}}\right) + d^0_{j,k}\log\frac{d^0_{j,k}}{d_{j,k}} - \left(d^0_{j,k} - d_{j,k}\right)\right], \tag{7}$$

with $2^J \geq 2^j \geq 2^{L+1}$. The existence of a minimum of $\mathscr{E}_{w_j,\varphi,\{a_{j,k}\},d^0_{j,k}}$ was proven in [9], where it was also shown that the Euler-Lagrange equations corresponding to the minimum are the following:

$$\frac{\partial\mathscr{E}_{w_j,\varphi,\{a_{j,k}\},d^0_{j,k}}}{\partial\{d_{j,k}\}}(d_{j,k}) = 0 \iff d_{j,k} = d^0_{j,k} + w_j\varphi'\left(\frac{a_{j,k}}{d_{j,k}}\right)\frac{a_{j,k}}{d_{j,k}}, \tag{8}$$

where φ' denotes the derivative of φ. In particular, when $\varphi \equiv \mathrm{id}$,

$$\frac{\partial\mathscr{E}_{w_j,\mathrm{id},\{a_{j,k}\},d^0_{j,k}}}{\partial\{d_{j,k}\}}(d_{j,k}) = 0 \iff d_{j,k} = d^0_{j,k} + w_j\frac{a_{j,k}}{d_{j,k}}, \tag{9}$$

and when $\varphi \equiv (\cdot)^\gamma$,

$$\frac{\partial\mathscr{E}_{w_j,\gamma,\{a_{j,k}\},d^0_{j,k}}}{\partial\{d_{j,k}\}}(d_{j,k}) = 0 \iff d_{j,k} = d^0_{j,k} + \gamma\omega_j\left(\frac{a_{j,k}}{d_{j,k}}\right)^\gamma. \tag{10}$$

Eqs. (9) and (10) are implicit equations that have to be solved using a numerical method. In [7, 9], Newton-Raphson's method has been used to find the zero of the function $F^{\varphi}(d_{j,k}) = d_{j,k} - d^0_{j,k} - \omega_j \varphi'\left(\frac{a_{j,k}}{d_{j,k}}\right)\frac{a_{j,k}}{d_{j,k}}$ by iteratively solving the equations

$$d^n_{j,k} = d^{n-1}_{j,k} - \frac{F^{\varphi}(d^{n-1}_{j,k})}{(F^{\varphi})'(d^{n-1}_{j,k})}, \quad n \geq 1. \tag{11}$$

Since the solution is not expected to differ too much from the original magnitude $d^0_{j,k}$, Newton-Raphson's algorithm is initialized with $d^0_{j,k}$.

The general scheme of the wavelet-based variational algorithm is the following:

1. We start by considering the approximation and detail coefficients of the coarsest level: $\{a_{J,k}, k \in \Omega\}$, and $\{d^{\ell}_{J,k}, k \in \Omega\}$, $\ell = H, V, D$;
2. The absolute value of the detail coefficients is modified according to the equations of eq. (11), by considering $a_{J,k}$ as fixed. This will implement local contrast enhancement in the wavelet domain at the scale 2^J;
3. We then pass to the scale 2^{J-1} and we compute the approximation coefficients by summing the approximation and detail coefficients just computed at the scale 2^J;
4. We fix these approximation coefficients and we repeat step 2 at the scale 2^{J-1};
5. We iterate this scheme until reaching the finest scale.

A general multiscale method for perceptually-inspired contrast enhancement of color images

In [9], the authors declare that the reason why wavelets have been chosen for their model is the fact that detail and approximation coefficients contain information about local contrast and local average, respectively, at different scales.

The fact that wavelets can be used to generate an orthonormal multilevel basis of the Hilbert space $L^2(\mathbb{R})$ or other specific properties of the wavelet theory have not been used in the model proposed by [9].

What is really important, is the possibility to have two finite sets, with the same cardinality, of bounded coefficients giving, at each fixed scale, a measure of local contrast and local average.

In this paper, we show that, for the aims of unsupervised contrast enhancement inspired by perceptual features, it is better to renounce to use of wavelets and build detail and approximation coefficients by using other strategies. The advantages will be both in the computational time and in the reduced artifacts affecting the results.

In the next subsection we give the details about the implementation of this idea with a Laplacian pyramid.

An implementation based on a Laplacian pyramid

The well known Laplacian pyramid decomposition [1] gives a multiscale representation of an image in which both approximation and detail coefficients are stored.

To build a Laplacian pyramid we first consider the original image and we operate the convolution with a Gaussian kernel, with an initial standard deviation $\sigma_0 = 0.25$, which will play an equivalent role as the spatial support of the mother wavelet. This results in a blurred (low pass filtered) version of the original image.

At each scale, we increase the blur by a factor $\sqrt{2}$. The blur for the j-th scale is therefore corresponding to the standard deviation $\sigma_j = \sigma_0 \sqrt{2}^j$. The iteration of blurring and subsampling generates a so-called Gaussian pyramid. In this way, each pixel contains a local average that corresponds to a pixel neighborhood on a lower level of the pyramid. This information is equivalent to that contained in the approximation coefficients of the wavelet framework.

The Gaussian pyramid is completed to a Laplacian one by saving the difference image of the blurred versions between each scale, these differences are the analog of the detail coefficients in the wavelet framework. As in the wavelet multiresolution analysis, only the coarsest scale is not a difference image. The Laplacian Pyramid bears this name because the difference of blurred versions is approximately equivalent to the convolution with the Laplacian of the Gaussian kernel.

We stop the decomposition when a threshold T based on the ratio between the minimum image dimension and σ_j is reached, that is if

$$\min(\text{width}, \text{height})/\sigma_j \leq T.$$

In our experiments, we use $T = 16$.

Having underlined the analogies between the Laplacian and the wavelet framework, it is immediate to see that formula (9) can be implemented by using for $d_{j,k}$ and $a_{j,k}$ the corresponding coefficients obtained via a Laplacian pyramid.

The results of this new algorithm, together with the comparison with the wavelet-based algorithm, will be discussed in the following section.

Tests

We have conducted intensive series of experiments to test the behavior of the new algorithm, in this section we report the results of these tests along with some example images.

First of all, let us consider the effect of the parameter α appearing in eq. (4). At least two options are available for α, for instance it can be: the per-channel average, i.e. the average of all the values of a fixed chromatic channel; the global image average, i.e. the average of all the values of all the channels. As it is logical to expect, if $\overline{a}_J$ is computed by averaging the approximation coefficients of the three chromatic channels, then it is more effective to remove a possible color cast affecting the image with respect to the option in which $\overline{a}_J$ is computed per-channel. Fig. 1 shows this effect.

Let us now consider the parameter w_j of eq. 5: of course, the bigger is w_j the higher is contrast enhancement, because the intensification of detail coefficients becomes stronger. Fig. 2 shows this effect.

Regarding the parameter γ: a small value of γ permits to use a bigger value of the weight w_j without introducing visible artifacts, while when γ grows, a more careful choice of w_j is required to avoid an excessive contrast intensification. Fig. 3 shows this effect.

Finally, we would like to show one of the main advantages of using a Laplacian pyramid instead of a wavelet basis. In Fig. 4 we compare the results of the two versions of the algorithm, showing that, with the same choice of parameters, the wavelet-based algorithm may show unwanted artifacts which depend on the shape of the mother wavelet chosen.

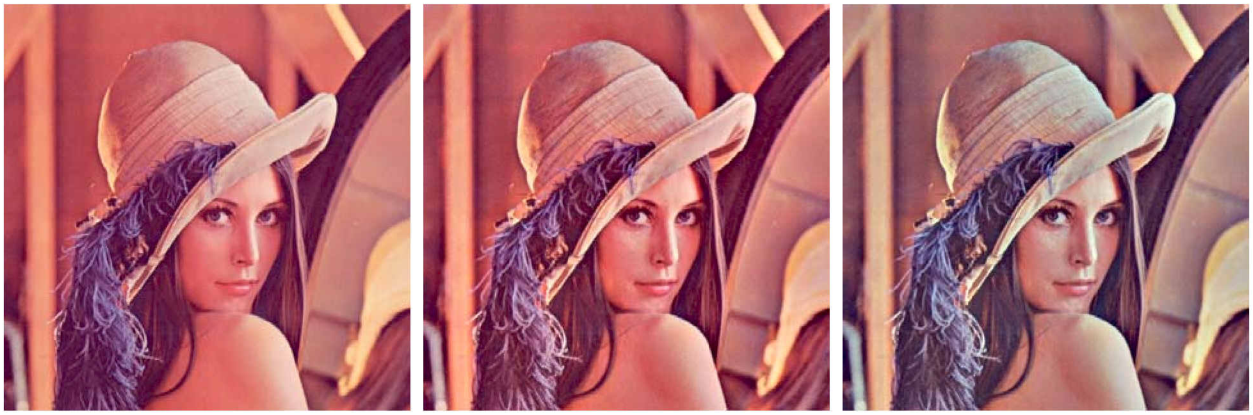

Figure 1. *Left: original image. Middle and Right: output images of the Laplacian pyramid version of the algorithm described by formula (10) with the following choice of parameters: $\gamma = 0.1$, $w_j = 0.05$ for all j, $\alpha = 0.5$, with the difference that the image in the center refers to a value of $\overline{a}_j$ computed per-channel, while the image on the right refers to $\overline{a}_j$ averaged along the three chromatic channels.*

Figure 2. *Left: original image. Middle and Right: output images of the Laplacian pyramid version of the algorithm described by formula (10) with the following choice of parameters: $\gamma = 0.1$, $\alpha = 0.5$, $w_j = 0.05$ for all j (center), $w_j = 0.05$ for all j, (right).*

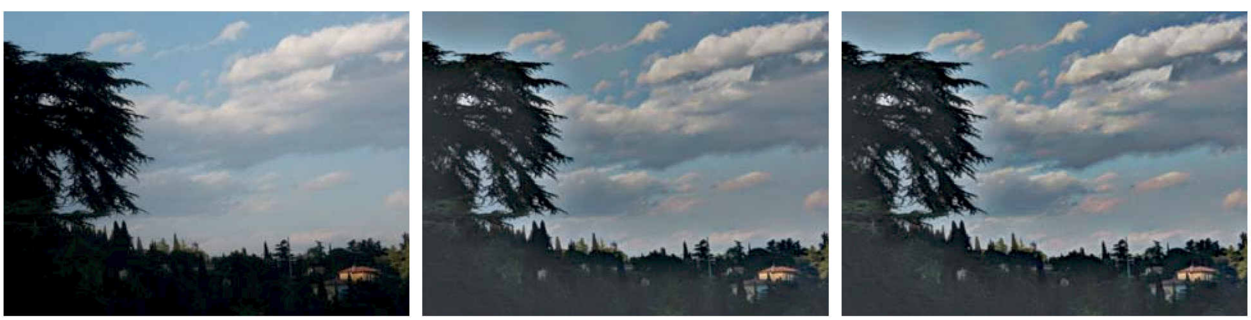

Figure 3. *Left: original image. Middle and Right: output images of the Laplacian pyramid version of the algorithm described by formula (10) with the following choice of parameters: $\alpha = 0.5$, $w_j = 0.05$ for all j, $\gamma = 0.1$ (center), $\alpha = 0.5$, $w_j = 0.005$ for all j, $\gamma = 1$ (right). Notice that we must considered weights ten times smaller when we set $\gamma = 1$, but even in that case we have some artifacts in the sky.*

Conclusions and perspectives

We have shown that the wavelet-based framework for perceptually-inspired contrast enhancement of color images proposed in [9], can be extended to any multiscale scheme in which both local approximation and detail coefficients are present. We have discussed, in particular, the case of the Laplacian pyramid, shown the effects of the parameters and compared its results with those of the original wavelet algorithm, showing that this new one is not affected by the typical wavelet-like artifacts.

References

[1] E. H. Adelson, C. H. Anderson, J. R. Bergen, P. J. Burt, and J. M. Ogden. Pyramid methods in image processing. *RCA engineer*, 29(6):33–41, 1984.

[2] M. Bertalmío, V. Caselles, and E. Provenzi. Issues about the retinex theory and contrast enhancement. *International Journal of Computer Vision*, 83:101–119, 2009.

[3] M. Bertalmío, V. Caselles, E. Provenzi, and A. Rizzi. Perceptual color correction through variational techniques. *IEEE Trans. on Image Processing*, 16:1058–1072, 2007.

Figure 4. *First row: original image (left), result of the Laplacian pyramid version of the algorithm described by formula (10) with parameters $\gamma = 0.1$, $\alpha = 0.5$, $w_j = 0.005$ for all j. Second and third row: results of the wavelet-based algorithm (with the same choice of parameters as above), with respect to the wavelets Sym2, Coif2 (second row, from left to right), CDF97 and Haar (third row, from left to right).*

[4] S. Ferradans, R. Palma-Amestoy, and E. Provenzi. An algorithmic analysis of variational models for perceptual local contrast enhancement. *Image Processing On Line*, 5(10):219–233, 2015.

[5] S. Mallat. *A Wavelet Tour of Signal Processing*. Academic Press, 3rd edition, 1999.

[6] R. Palma-Amestoy, E. Provenzi, M. Bertalmío, and V. Caselles. A perceptually inspired variational framework for color enhancement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(3):458–474, 2009.

[7] E Provenzi. Boosting the stability of wavelet-based contrast enhancement of color images through gamma transformations. *Journal of Modern Optics*, 60:1145–1150, 2013.

[8] E. Provenzi. *Computational Color Science: Variational Retinex-like Methods*. John Wiley & Sons, 2017.

[9] E. Provenzi and V. Caselles. A wavelet perspective on variational perceptually-inspired color enhancement. *International Journal of Computer Vision*, 106:153–171, January 2014.

[10] E. Provenzi, L. De Carli, A. Rizzi, and D. Marini. Mathematical definition and analysis of the retinex algorithm. *Journal of the Optical Society of America A*, 22(12):2613–2621, December 2005.

[11] E. Provenzi, M. Fierro, A. Rizzi, L. De Carli, D. Gadia, and D. Marini. Random spray retinex: A new retinex implementation to investigate the local properties of the model. *IEEE Transactions on Image Processing*, 16:162–171, January 2007.

[12] E. Provenzi, C. Gatta, M. Fierro, and A. Rizzi. Spatially variant white patch and gray world method for color image enhancement driven by local contrast. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 30:1757–1770, 2008.

[13] A. Rizzi, C. Gatta, and D. Marini. A new algorithm for unsupervised global and local color correction. *Pattern Recognition Letters*, 24:1663–1677, 2003.