Space Sensitive Color Gamut Mapping: A Variational Approach

Ron Kimmel[1], Doron Shaked[2], Michael Elad[3], Irwin Sobel[4]
HP Laboratories
HPL-2000-50 (R.1)
March 15th , 2002*

gamut mapping, retinex algorithm, calculus of variations, robust reconstruction

Gamut mapping is a method to modify a representation of a color image to fit into a constrained color space of a given rendering medium. A laser-jet color printer that attempts to reproduce a color image on a regular paper would have to map the photographed picture colors in a given color range, also known as the image 'color gamut', into the given printer/page color gamut. Most of the classical gamut mapping methods involve a pixel by pixel mapping and ignore the spatial color configuration. Recently proposed spatial dependent approaches for gamut mapping are either based on heuristic assumptions or involve a high computational cost. Note that spatially varying gamut mapping is inherently image dependent.

We present a new variational approach for space dependent gamut mapping. The approach presents a new measure for the problem, and is closely related to a recent variational framework for Retinex. We link our method to recent measures that attempt to couple spectral and spatial perceptual measures. We show that the solution to our formulation of the problem is unique if the gamut of the target device is convex. A quadratic programming efficient numerical solution is proposed, with real-time promising results.

* Internal Accession Date Only Approved for External Publication
[1] The Computer Science Department, Technion, Haifa 32000, Israel
[2] HP Labs Israel, Technion City, Haifa 32000, Israel
[3] The Computer Science Department, Stanford, Palo Alto, CA, 94306
[4] HP Labs, Palo Alto, 1501 Page Mill Rd, Palo Alto, CA 94304

1 Introduction

Gamut mapping is a method to modify a representation of a color image to fit into a constrained color space of a given rendering medium. A laser-jet color printer that attempts to reproduce a color image on a regular paper would have to map the photographed picture colors in a given color range, also known as the image 'color gamut', into the given printer/page color gamut. Most of the classical gamut mapping methods involve a pixel by pixel mapping (usually a pre-defined look-up table) and ignore the spatial color configuration. Only recently, spatial dependent approaches where proposed for gamut mapping [6, 7]. However these solutions are either based on heuristic assumptions or involve a high computational cost.

We present a new variational approach for space dependent gamut mapping. The approach presents a new measure for the problem, and is closely related to a recent variational framework for the Retinex [4, 5]. We link our method to recent measures that attempt to couple spectral and spatial perceptual measures [7, 9]. We show that the solution to our formulation of the problem is unique if the gamut of the target device is convex. A quadratic programming efficient numerical solution is proposed, with real-time promising results. Note that spatially varying gamut mapping is inherently image dependent.

The structure of the paper is as follows. Section 2 reviews recent previous work on space sensitive gamut mapping and perceptual measures for the spectral-spatial case. Next, Section 3 introduces the proposed framework. We start from the functional definition, derive its Euler Lagrange as a gradient descent process, describe the numerical approximation, comment on uniqueness and convergence, and the relation to the Retinex problem. We conclude with experimental results comparing the proposed method to alternative algorithms on a set of images.

2 Previous Work

In his patent application [6], McCann suggests to preserve spatial gradients in all scales while applying a gamut mapping procedure. The basic idea is to preserve the magnitude of the gradients in the original image, while projecting onto the target gamut as a constraint. The multi-scale property is achieved by sampling the image around each pixel with exponentially increasing sampling intervals while the sampling is done along the vertical and horizontal directions. McCann's method preserves image gradients. Nevertheless, a better approach would be to start from an objective measure (a functional) for gradient preservation and other imaging goals. A sound mathematical foundation is bound to give a good understanding of the problem and the inherent trade-offs, and consequently improve practical solutions.

A simple spatial-spectral measure for human color perception was proposed by Zhang and Wandell [9]. The 'S-CIELAB' defines a spatial-spectral measure for human color perception by a composition of spatial band-pass linear filters in the opponent color space followed by the CIELAB Euclidean perceptual color measure [9]. We latter link between 'S-CIELAB' and our proposed method.

In [7], Nakauchi, Hatanaka, and Usui, modulate an L_2 measure for image difference by *human contrast sensitivity* functions. The authors use a model in which the contrast sensitivity function is a linear combination of three spatial band-pass filters H_1, H_2, H_3 given in the spatial-frequency domain (or h_1, h_2, h_3, as their corresponding spatial filters), see Figure 1.

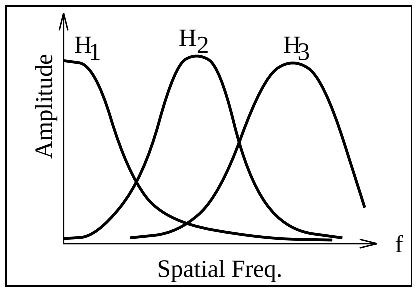

Figure 1: Qualitative description of filters modeling the *human contrast sensitivity* functions, in the spatial frequency domain

For gamut mapping of the image u_0 in the CIELAB space, Nakauchi et al. minimize the functional

$$E(u^L, u^a, u^b) = \sum_{i=1}^{3} \sum_{c \in \{L,a,b\}} \int_\Omega \left(h_i^c * (u^c - u_0^c)\right)^2 d\Omega, \tag{1}$$

subject to $\{u^L, u^a, u^b\} \in \mathcal{G}$. Where h_i^c is the filter corresponding to the spectral channel $c \in \{L, a, b\}$ and the $i \in \{1, 2, 3\}$ 'contrast sensitivity' mod. Ω is the image domain, and $\mathcal{G}$ the target gamut. Note that a total of nine filters are involved, three for each spectral channel and a total of three spectral channels.

The filters H_i^c are modeled by shifted Gaussians. H_1^c is not shifted, and thus, h_1^c is also Gaussians, while h_2^c, and h_3^c are a Gaussian modulated by two sinus functions with different periods. A graphical analysis of h_2^c, and h_3^c, as in Figure 2, argues that they approximate the derivative operator at different scales. Denote these two gradient approximation operators by $\nabla_{\sigma_1}^c$ and $\nabla_{\sigma_2}^c$. Note that any band pass filter can be considered as a version of a derivative operator. Furthermore, one possible extension of the 1D derivative to 2D is the gradient.

Thus, we maintain that minimization of Nakauchi's functional (1) is similar to minimizing the following functional for each channel separately

$$\int_\Omega |h_1^c * (u - u_0)|^2 + |\nabla_{\sigma_1}^c (u - u_0)|^2 + |\nabla_{\sigma_2}^c (u - u_0)|^2 d\Omega. \tag{2}$$

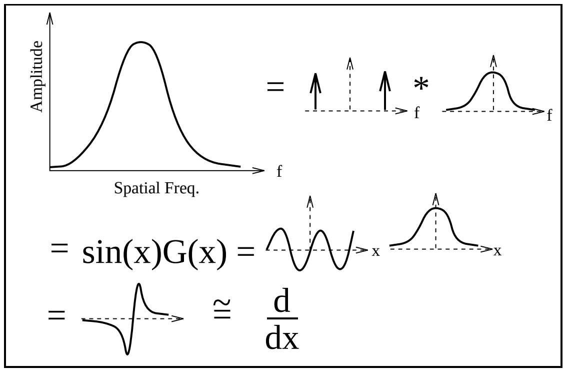

Figure 2: Shifted Gaussian is roughly a derivative operation

Roughly speaking, the first term corresponds to the S-CIELAB perceptual measure, while the next two terms capture the need for matching the image variations at two selected scales that were determined by human perception models. One technical difficulty of the spatial filters corresponding to (1) is their large numerical support, which is probably the reason for the slow implementation reported in [7]. Next, we show an alternative view of the problem with an efficient numerical solution.

3 The Proposed Variational Model

A good measure of image deviation captures the perceptual difference between the initial, u_0, and final, u, images. This is modeled by

$$\mathcal{D} = g * (u - u_0). \tag{3}$$

where g is, say, a normalized Gaussian kernel with zero mean and a small variance σ. This model is good for small deviations. However, for large deviations it should be elaborated to account for possible perceptual *feature* differences, which may be modeled by the difference of gradients, which due to linearity, turns out to be the gradient of (3)

$$\nabla\mathcal{D} = \nabla\left[g * (u - u_0)\right] = g * (\nabla u - \nabla u_0). \tag{4}$$

The proposed measure yields the functional

$$E(u) = \int_\Omega \left(\mathcal{D}^2 + \alpha|\nabla\mathcal{D}|^2\right) d\Omega. \tag{5}$$

which should be minimized subject to $u \in \mathcal{G}$. Note the similarity between the above and (2), which is a propper measure for image procesing. Note also that the above is a Sobolev space norm.

Taking the first variation of (5) w.r.t. u we get the EL equation

$$\frac{\delta E(u)}{\delta u} = g * (\alpha \text{div} (\nabla g * (u - u_0)) - g * (u - u_0)) = 0, \tag{6}$$

see Appendix A for this derivation. Reformulating the EL as a gradient descent flow for u, we get the following minimization scheme

$$\begin{cases} \frac{du}{dt} = \alpha g * \Delta \mathcal{D} - g * \mathcal{D}, \\ \text{s.t.} \quad u \in \mathcal{G} \end{cases} \tag{7}$$

3.1 Properties of the Model

The proposed functional and the resulting minimization scheme are both Euclidean invariant in the image plane. They are thus both translation and rotation invariant. As the parameter α goes to zero we approximate the S-CIELAB model, while for effective α we have a proper extension to the perceptual measures proposed in [7], with an efficient numerical implementation.

3.2 Numerical Implementation

Recall, that we added an artificial time parameter t to the image $u(x, y)$, that now reads $u(x, y; t)$. Let us discretize the EL gradient descent equation, by first taking a simple forward explicit approximation for the t derivative,

$$\frac{u^{n+1} - u^n}{\tau} = \alpha g * \Delta \mathcal{D} - g * \mathcal{D},$$

where $\tau = dt$ and $u^n(x, y) \approx u(x, y; n\tau)$.

Next, we deal with the space derivatives. Let $u^n_{i,j} \approx u(ih, jh; n\tau)$, where we assume uniform spatial spacings in the x and y directions of size h. We use central derivatives in space,

$$\begin{aligned} u_{xx} &\approx D_{xx} u &\equiv \frac{u_{i+1} - 2u_i + u_{i-1}}{h^2} \\ u_x &\approx D_x u &\equiv \frac{u_{i+1} - u_{i-1}}{2h}, \end{aligned}$$

and the same for the y direction. We also use the relation $g * D_{xx}(g * u) = g_x * g_x * u$, and compute the kernels $D_2 = g_x * g_x + g_y * g_y = D_x g * D_x g + D_y g * D_y g$. The explicit approximation reads

$$\begin{aligned} \tilde{\mathcal{D}}^n &= g * g * (u^n - u_0) \\ \mathcal{L}^n &= D_2 * (u^n - u_0) \\ u^{n+1}_{ij} &= u^n_{ij} + \tau \left(\alpha \mathcal{L}^n_{ij} - \tilde{\mathcal{D}}^n_{ij} \right). \end{aligned}$$

subject to the constraint $u_{ij}^n \in \mathcal{G}$.

In order to speed up convergence, we used a standard coarse to fine pyramidal approach. A full multi-grid method is another possibility, in which the joint projection introduces some interesting problems, that will be explored elsewhere.

3.3 Uniqueness and Convergence

The proposed functional has a Quadratic Programming (QP) form, since the penalty term is quadratic and the constraint is linear. If the set $\mathcal{G}$ is convex, the overall problem is convex if and only if the Hessian of the functional is positive definite [2]. In such a case, there is a unique local minimum which is obviously also the global solution to the problem. In our case, the Hessian is given by $g * (1 - \alpha\Delta) * g$, which is indeed positive definite for all $\alpha > 0$. Thus, for a convex target gamut $\mathcal{G}$, there exists a unique solution.

3.4 Relation to Retinex

The gamut mapping problem is related to the Retinex problem of illumination compensation and dynamic range compression. The basic Retinex problem is: How to estimate the reflectance image from the given acquired image ? A reasonable optical model of the acquired image S asserts that it is a multiplication of the reflectance R and the illumination L imagesr. Where the reflectance image is a hypothetic image that would have been measured if every visible surface would have been illuminated by a unit valued white illumination source, and the illumination image is the actual illumination shaded on surfaces in the scene. In the log domain we get

$$s = r + l$$

where s, r, and l are the respective logarithms of S, R, and L. Since we know that the surface patches can not reflect more light than has been shaded on them $R < 1 \implies r < 0$. Thus, we want an image $r < 0$, which is perceptually similar to s. For the Retinex we have an additional physically motivated constraint, namely, that the illumination image $l = s - r$ is smooth, i.e. the gradient $|\nabla l| = |\nabla (r - s)|$ is small. But this is just another way to say that the features of r are similar to those of s, since we do not assume that the illumination created perceptual features in s. In the gamut mapping problem we have an image u_0, and we want to estimate an image $u \in \mathcal{G}$ which is not only perceptually similar to u_0, but also has similar perceptual features as u_0.

3.5 Robust Version of the Proposed Algorithm

The proposed penalty function as shown in Equation 5 tends to create halows in the resulting image. Figure 3 explains the origin of those halows through a one dimensional example. In Figure 3 we see a signal which is outside of the gamut (marked by dotted lines). Projecting the signal onto the gamut will result in a constant value and loss of all detail. The dashed line represents the result of scaling the signal into the allowed range. All the details are preserved, but with a smaller contrast. As opposed to these point operations, our space dependent approach yields a signal which preserves the details with high contrast (the solid line). However, near the strong edges we get halows, which means that near the edge there is a slow transition from low to high values.

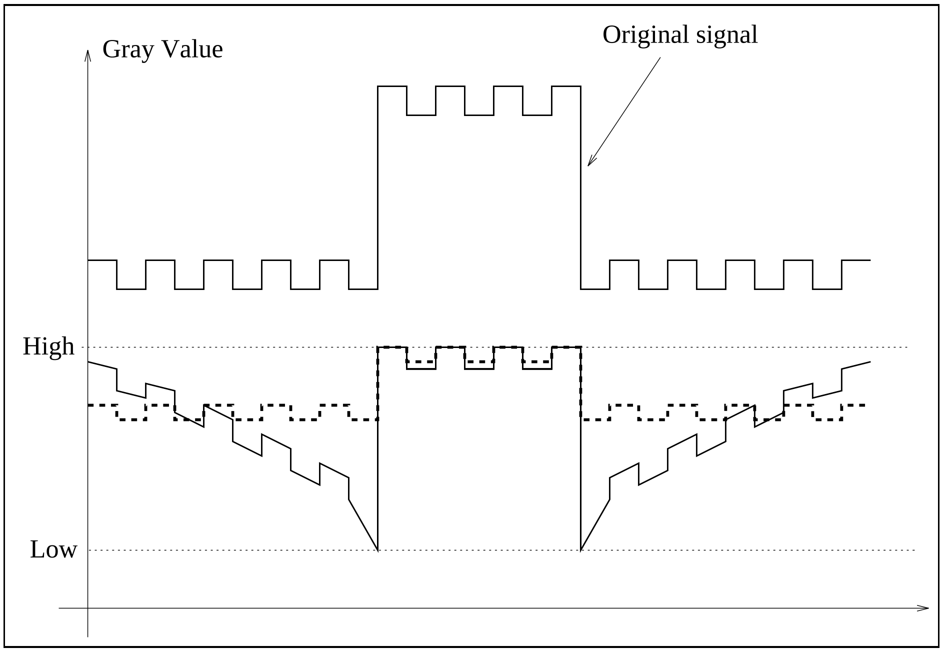

Figure 3: A 1D example of the algorithm's behavior and the creation of halows

In order to avoid this phenomena, we have to modify our penalty term, and use robust estimation tools. The original penalty term in (5) should be replaced by

$$E(u) = \int_{\Omega} \rho_1(\mathcal{D}) + \alpha \rho_2(|\nabla \mathcal{D}|) d\Omega. \tag{8}$$

which for $\rho_1(x) = \rho_2(x) = x^2$ coincides with (5). If the function $\rho(x)$ grows slower than x^2 as $x \to \infty$, we get an improved behavior near strong edges. Good candidates for $\rho(x)$ are $\rho(x) = |x|$ or $\rho(x) = \sqrt{1 + x^2}$.

A different simpler (linear) approach with similar robust behavior is to solve the original problem (5) twice, with two different values of α. We denote the solution with a small α as u_{small} and the one which corresponds to the high value of α as u_{high}. The solution u_{small} has smaller contrast at areas with small details, yet has almost no halows. On the other hand, u_{high} preserves the small details, but at the expense of strong halow effects. Therefore, by averaging these two results in a spatially adaptive way, we can enjoy both worlds. The proposed solution is therefore

$$u_{final}[k,j] = w[k,j]u_{small}[k,j] + (1 - w[k,j])\, u_{high}[k,j]$$

The weight $w[k,j]$ should be close to one near strong edges, and close to zero in relatively smooth regions. In our experiments we used

$$w[k,j] = \frac{1}{1 + \beta |\nabla g * u_0|^2}$$

and achieved resemblance to robust estimation.

Halow problems have been recently dealt with in relation to Dynamic Range Compression. Solutions proposed included anisotropic diffusion [8] and robust filtering [3] (refer to [1] for connection between these two approaches). The solutions proposed in this section are solutions to the same halow problem, and have been proven efficient also for the variational Retinex [5].

4 Results

Figures 4 and 5 present the result of the proposed measure compared to the regular L_2 norm. In this example we used two resolution levels with four iterations at each resolution, $\alpha = 10, dt = 0.0011$, $\sigma = 1.1$, and the support of the Gaussian kernel is set to 15×15 pixels. The minimization in this example was applied in the RGB, space. We defined the target gamut such that the RGB channels are restricted to the range $[40, 100]$ instead of the gamut range $[0, 255]$ of the original image. We also tested the effect of the SCIELAB measure in the opponent color space, without any dramatic effect on the final result.

We turn now to present the robust gamut mapping results. The applied algorithm is the shortcut method of adaptive weighting two regular results, obtained with different values of α. Figure 6 presents an original image, and the two solutions obtained by the regular variational penalty function with $\alpha = 1$ and $\alpha = 40$. The limited gamut in this case is as before, namely, R, G and B values in the range $[40, 100]$. Figure 7 show the weight image as computed by the proposed formula with $\beta = 0.005$, and the weighted average result. As can be seen, halows are suppressed in the final result, while preserving the details of the input image.

5 Concluding Remarks

We presented a variational formulation for the gamut mapping as a QP problem. A simple functional that measures both the image difference and its derivatives was shown to be analog to perceptual difference measures. Actually, this is a similarity measure in Sobolev space in which the proximity of the derivatives capture the small scale and the detailed information

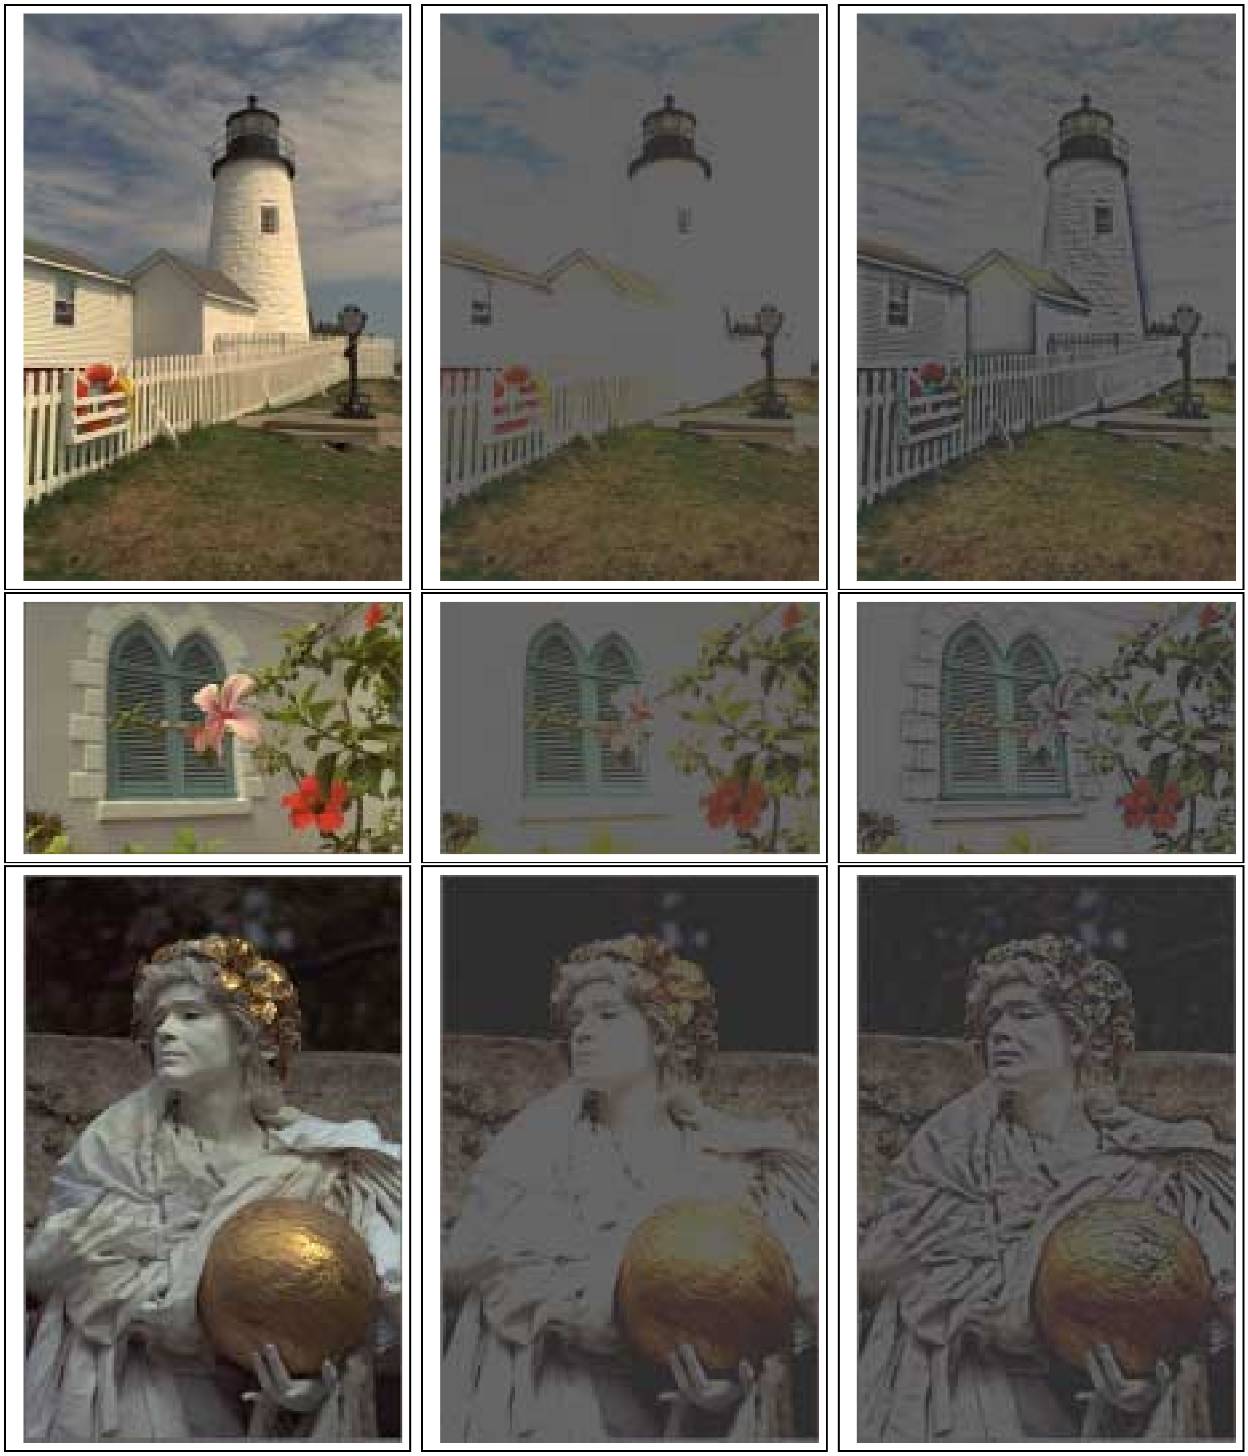

Figure 4: Left to right: Original image, gamut mapping by truncation (minimization of the L_2 norm), and the result of the proposed scheme.

Figure 5: Left to right: Original image, gamut mapping by truncation (minimization of the L_2 norm), and the result of the proposed scheme.

Figure 6: Left to right: Original lighthouse image, gamut mapping by the original penalty function with $\alpha = 1$, and the same with $\alpha = 40$.

Figure 7: Left to right: the weight image (white=1, black=0), and the final weighted average image.

of the difference between the images. We linked our results to previous methods including solutions to the Retinex problem, and presented an efficient numerical multi-resolution algorithm for its solution, which can be used for image reproduction subject to convex constraints with a unique solution.

Acknowledgments

The authors would like to thank Nathan Moroney from HP Labs Palo Alto for pointing out the important reference [7], and to Avi Levi from HP Labs Israel for intriguing discussions on Sobolev Spaces and the Gamut problem during the summer of 1999.

Appendix A

Let us explore the effect of a convolution operation within the functional on the Euler Lagrange equation.

First, by linearity, it is simple to show that for a symmetric kernel g we have $\frac{d}{dx}(g*u) = (\frac{d}{dx}g)*u = g*(\frac{d}{dx}u)$, or in short hand notations, $\frac{d}{dx}(g*u) = g_x * u$. Next, given the general functional of the form

$$E(u) = \int F\left(\frac{d}{dx}(g*u)\right) dx,$$

we set $\tilde{u}(x) = u(x) + \epsilon\eta(x)$, and calculate

$$\begin{aligned}
\frac{dE(\tilde{u})}{d\epsilon} &= \frac{d}{d\epsilon}\int F\left(\frac{d}{dx}(g*\tilde{u})\right) dx \\
&= \int \frac{d}{d\epsilon} F\left(\frac{d}{dx}\int g(x-\hat{x})\tilde{u}(\hat{x})d\hat{x}\right) dx \\
&= \int F'\left(\frac{d}{dx}(g*\tilde{u})\right)\frac{d}{d\epsilon}\left(\frac{d}{dx}\int g(x-\hat{x})\tilde{u}(\hat{x})d\hat{x}\right) dx \\
&= \int F'(g_x*\tilde{u})\frac{d}{dx}\left(\int g(x-\hat{x})\frac{d}{d\epsilon}\tilde{u}(\hat{x})d\hat{x}\right) dx
\end{aligned}$$

Using integration by parts and vanishing boundary values, i.e. $\int uv' = -\int vu'$, we get

$$\begin{aligned}
\frac{dE(\tilde{u})}{d\epsilon} &= -\int (g*\eta)\frac{d}{dx}F'(g_x*\tilde{u})dx \\
&= -\int\int g(x-\hat{x})\eta(\hat{x})d\hat{x}\frac{d}{dx}F'(g_x*\tilde{u})dx \\
&= -\int \eta(\hat{x})\int g(x-\hat{x})\frac{d}{dx}F'(g_x*\tilde{u})dx d\hat{x} \\
&= -\int \eta\left[g*\left(\frac{d}{dx}F'(g_x*\tilde{u})\right)\right] d\hat{x}.
\end{aligned}$$

The extremum condition is checked in the limit, as $\epsilon \to 0$, such that $dE/d\epsilon = 0$ for all $\eta(x)$. It is given by the Euler Lagrange equation

$$g * \frac{d}{dx} F'(g_x * u) = 0,$$

or equivalently $g_x * F'(g_x * u) = 0$.

For example, for the functional

$$E(u) = \int \left(\frac{d}{dx} \left(g * (u - u_0) \right) \right)^2 dx,$$

the EL is given by

$$g * \frac{d}{dx} \left(\frac{d}{dx} (g * (u - u_0)) \right) = 0.$$

or equivalently $g * g_{xx} * (u - u_0) = 0$.

6 References

[1] D. Barash, "Bilateral Filtering and Anisotropic Diffusion: Towards a Unified Viewpoint", *Hewlett-Packard Laboratories Technical Report 18-2000*, February 2000.

[2] D. P. Bertsekas, Non-linear Programming, Athena Scientific, 1995.

[3] J. M. DiCarlo, and B. A. Wandell, "Rendering High Dynamic Range Images", *Proc. of Electronic Imaging*, January 2000.

[4] M. Elad, R. Kimmel, D. Shaked, and R. Keshet, "Reduced Complexity Retinex Algorithm via the Variational Approach", *Hewlett-Packard Laboratories Technical Report19-2000*, February 2000.

[5] R. Kimmel, M. Elad, D. Shaked, R. Keshet, and I. Sobel, "A Variational Framework for Retinex", *Hewlett-Packard Laboratories Technical Report 99-151*, November 1999.

[6] J. J. McCann, Patent application, January 1999

[7] S. Nakauchi, S. Hatanaka, and S. Usui, "Color Gamut Mapping Based on a Perceptual Image Difference Measure" *Color research and application*, Vol. 24, pp. 280–291, 1999.

[8] J. Tumblin, and G. Turk, "LCIS: A Boundary Hierarchy for Detail-Preserving Contrast Reduction", *SIGGRAPH*, 1999.

[9] X. Zhang and B. A. Wandell, "A Spatial Extension of CIELAB for Digital Color Image Reproduction" *Proc. of the SID Symposium*, Vol. 27, pp. 731-734, 1996