

# Space-Dependent Color Gamut Mapping: A Variational Approach

Ron Kimmel, Doron Shaked, Michael Elad, and Irwin Sobel

**Abstract**—Gamut mapping deals with the need to adjust a color image to fit into the constrained color gamut of a given rendering medium. A typical use for this tool is the reproduction of a color image prior to its printing, such that it exploits best the given printer/medium color gamut, namely the colors the printer can produce on the given medium. 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. In this paper, we present a new variational approach for space-dependent gamut mapping. Our treatment starts with the presentation of a new measure for the problem, closely related to a recent measure proposed for Retinex. We also link our method to recent measures that attempt to couple spectral and spatial perceptual measures. It is shown that the gamut mapping problem leads to a quadratic programming formulation, guaranteed to have a unique solution if the gamut of the target device is convex. An efficient numerical solution is proposed with promising results.

## I. INTRODUCTION

THE term “color gamut” stands for the span of all possible colors of a given image or device. For an image, the color gamut is simply the set of all the colors found in it. For output devices, such as printers or screens, the color gamut is the set of colors the given device can render. A similar definition exists for input devices.

When a color image is to be rendered on a printer or a screen, “gamut mapping” is typically required. This process deals with the need to adjust the colors of the input image such that they fit into the constrained color gamut of the output device. The gamut mapping problem is, thus, a fundamental one in any transfer of color images from input to output devices. Clearly, in designing such mappings, one should be interested in transforms that, while fulfilling the basic desire to match the gamuts, preserve the original details as much as possible.

Most of the classical gamut mapping methods involve a pixel-by-pixel mapping (usually a predefined lookup table) and ignore the spatial color configuration in the rendered image. Only recently, spatial-dependent approaches were proposed [1], [13], [14]. However, these solutions are either based on heuristic assumptions or involve a high computational cost.

Manuscript received July 17, 2003; revised June 14, 2004. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Zhigang Fan.

R. Kimmel and M. Elad are with The Computer Science Department, The Technion—Israel Institute of Technology, Haifa 32000, Israel.

D. Shaked is with Hewlett-Packard Laboratories Israel, Technion City, Haifa 32000, Israel.

I. Sobel is also with the Hewlett-Packard Laboratories, Palo Alto, CA 94304 USA.

Digital Object Identifier 10.1109/TIP.2005.847299

One of the fundamental motivations of spatial gamut mapping is the need to preserve the edge between two out-of-gamut colors, which would otherwise map individually to the same in-gamut color. Nonspatial gamut mapping approaches such as projection or shrinkage would either map the two colors to the same in-gamut color (projection), thereby eliminating the edge, or map them to different colors (shrinkage), thereby distorting at least one of the colors in cases where it appears with no similar color in its spatial vicinity.

Spatial gamut mapping approaches will optimally map the two out-of-gamut colors to the same in-gamut color, when they appear separate in the image, and to different colors when the two input colors share an edge. In some algorithms, these requirements constitute a conflict. For example, suppose the two colors mentioned above share an edge, but extend each in a large constant tone patch to cover a large part of the image. What would be the best mapping for a pixel located far from the edge between the two colors. Should it be mapped as if it were close to the edge or as if it were disconnected from the edge?

In this paper, we present a new variational approach for space-dependent gamut mapping. Due to its adaptation to the spatial color configuration behavior, the proposed method is also image dependent by definition. Our analysis starts by presenting a new metric between two candidate representations of the same color image. This measure is shown to be closely related to a recent measure based on a variational framework proposed for the Retinex problem [6], [8], [9]. The proposed measure is also linked to recent work that coupled spectral and spatial perceptual distances into one tool [14], [19].

Using the proposed distance measure, it is shown that the gamut mapping problem translates to a quadratic programming optimization form. This problem becomes convex if the gamut of the target device is convex. For such devices the proposed method is, thus, guaranteed to yield a unique optimal solution. A highly efficient pyramidal (simplified multigrid) numerical algorithm to find this solution is proposed with promising results.

Gamut mapping is a practical problem, with many technical aspects and difficulties. In this paper, we chose to restrict our scope to a mathematical framework addressing spatial considerations in gamut mapping. In doing so, we deliberately deemphasized other important issues such as color fidelity and projection type. The main role of this paper is, therefore, showing that the spatial considerations can be taken care of by a variational approach. This effectively minimizes a metric for the distance between images, taking spatial considerations into effect. We believe that this concept is flexible, and can be merged with an arbitrary projection and color-fidelity measure, all in the desire to design a practical end-to-end gamut mapping algorithm.

The structure of the paper is as follows. Section II reviews recent work on space-dependent gamut mapping and perceptual measures for the spectral-spatial case. Next, in Section III, we introduce the proposed framework. We start from the functional definition, derive its Euler–Lagrange (EL) as a gradient descent process, describe the numerical approximation, comment on uniqueness and convergence, make relation to the Retinex problem, and, finally, suggest methods to robustify the algorithm to better treat halos. In Section IV, we present experimental results comparing the proposed method to non-spatial gamut mapping on a set of images. We conclude and summarize in Section V.

## II. PREVIOUS WORK

Among the various existing methods for the design of gamut mapping, the method proposed by McCann [13] is the closest in spirit to ours. McCann suggests to preserve spatial gradients in all scales while applying a gamut mapping procedure. The basic idea is the preservation of the gradients magnitude as in the original image, while projecting onto the target gamut as a constraint. The multiscale property is achieved by sampling the image around each pixel with exponentially decreasing sampling intervals. Sampling in [13] is done along the vertical and horizontal directions. Generally speaking, McCann’s algorithm exhibits good performance, both in terms of speed and output quality. An annoying halo artifact may sometimes appear near sharp edges.

A different method to pronounce gradients in the gamut mapping process is reported by Bala, deQueiroz, Eschbach, and Wu [1]. They suggest to use nonspatial projection gamut mapping and compute the difference between the luminance channels of the result and the original image. A high-pass filtered version of this difference is then added to the luminance channel of the mapped result, this way pronouncing edges. Finally, the edge enhanced image needs to be mapped again to the new gamut.

The above papers take a different approach toward the basic problem of spatial gamut mapping. The former algorithm by McCann [13] arbitrates the edge preservation and color fidelity requirements indirectly via heuristic choice of parameters. A wrong choice may be reflected in halos near sharp edges. Bala *et al.* [1] applies only one cycle of this process, imposing the spatial versus gamut conflicting requirements. Out-of-gamut color edges affect the output only to the extent of the support of the high pass filter. While very effective, this approach is analogous to using edge sharpening to solve contrast problems.

Whereas the above methods are indeed successful in preserving image gradients, they are based mainly on heuristics and knowhow. As a consequence, it is hard to discuss their implied metrics, or in other words, what it is that makes two images different: gradients versus tone differences and those in the luminance versus chrominance channels. Furthermore, it is hard to analyze general algorithm properties such as convergence to optimality, choice of parameters, possible improvements, and more.

The gamut mapping method proposed in this paper is also directed toward preserving gradients. However, as this new method starts from an objective measure (a functional), it builds

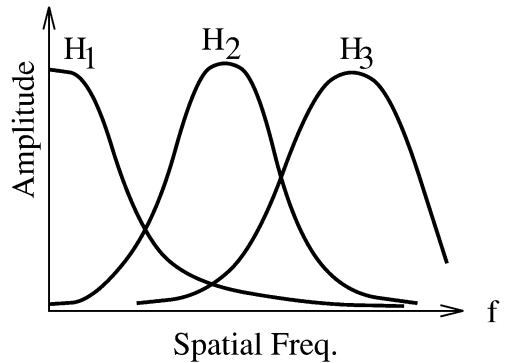


Fig. 1. Qualitative description of filters modeling the *human contrast sensitivity* functions, in the spatial frequency domain, as in [14].

a sound mathematical foundation. This approach gives a good understanding of the problem and its inherent tradeoffs and, consequently, leads to an improved practical solution.

While working on this problem, it became clear to us that the key for obtaining a clear formulation for the gamut mapping task is the construction of a mathematical model for proximity between two different representations of the same color image. A marked progress toward this end was accomplished by Zhang and Wandell [19]. They suggest a simple spatial-spectral measure for human color perception, called the S-CIELAB. This measure 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 [19]. We latter link S-CIELAB and our method.

A similar effort to define such a distance measure is found in [14]. In this work, 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 Fig. 1).

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

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

subject to  $\{u^L, u^a, u^b\} \in \mathcal{G}$ . Here, the symbol  $*$  denotes convolution, and  $h_i^c$  is the filter corresponding to the spectral channel  $c \in \{L, a, b\}$  and the  $i \in \{1, 2, 3\}$  contrast sensitivity mode.  $\Omega$  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 Gaussian functions.  $H_1^c$  is not shifted, and, thus,  $h_1^c$  is simple Gaussian, whereas  $h_2^c$  and  $h_3^c$  are Gaussians modulated by two sine functions with different frequencies. One may argue that these shifted smoothers approximate the smoothed derivative operator at different scales. Thus, we maintain that minimization of Nakauchi’s

functional (1) is similar to minimizing the following functional for each channel separately

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

where  $\nabla_{\sigma_1}^c$  and  $\nabla_{\sigma_2}^c$  are two smoothed derivatives, replacing the filtering done by  $h_2^c$  and  $h_3^c$ , respectively. 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 support, which is probably the reason for the slow implementation reported in [14]. Next, we show an alternative view of the problem with an efficient numerical solution.

### III. PROPOSED SOLUTION

#### A. Modeling the Problem

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

$$\mathcal{D}^c = g * (u^c - u_0^c) \quad (3)$$

where  $c \in \{L, a, b\}$ ,  $*$  denotes convolution, and  $g$  is a normalized Gaussian kernel with zero mean and a small variance  $\sigma$ . This model is good for small deviations. However, as deviations become larger, it should be modified 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)

$$\nabla \mathcal{D}^c = \nabla [g * (u^c - u_0^c)] = g * (\nabla u^c - \nabla u_0^c). \quad (4)$$

The proposed proximity measure between  $u^c$  and  $u_0^c$  yields the functional

$$E(u^L, u^a, u^b) = \frac{1}{2} \sum_{c \in \{L, a, b\}} \int_{\Omega} (\|\mathcal{D}^c\|_C + \alpha \|\nabla \mathcal{D}^c\|_G) d\Omega \quad (5)$$

which should be minimized subject to  $u \in \mathcal{G}$ . Note that we deliberately leave the norms involved ( $C$  and  $G$ ) as general and undefined yet. The choices for the norms lead essentially to the choice of projection mentioned earlier.

The proposed proximity measure is similar to (2), the proximity measure implied in [14]. It is also a Sobolev space norm [12]. The expression in (3) reflects our expectation to have the same colors in the output and input. Indeed, if  $\alpha = 0$ , the optimal solution to (5) is nonspatial projection. The projection type used is determined by the choice of the norm  $\|\cdot\|_C$ . For example, if  $\|\cdot\|_C$  weights the luminance channel much less than the chrominance channel, the projection will preserve chrominance. On the other hand, if all components are weighted equally, the projection is orthogonal in the formulation's color space.

The metric implied in (5) reflects our expectation to have the same color gradients in the output and input. Indeed if  $\alpha \rightarrow \infty$  the optimal solution to (5) will tend to move the image in the color space and fit it into  $\mathcal{G}$ . Assuming that this is impossible, some edges (primarily the smaller ones) will be shrunken in order to enable the desired fit. The projection type used for the resultant shrinkage can be determined by the choice of the norm  $\|\cdot\|_G$ . For example, if  $\|\cdot\|_G$  weights the luminance channel much less than the chrominance channel, the projection will preserve chrominance (much like what happens in [1] where the edge enhancement is restricted to the luminance channel). On the other hand, if all components are weighted equally, the projection is orthogonal in the formulation's color space, and, thus, edges are preserved in both luminance and chrominance.

We are aware that, in practice, one may want to work with a different metric.<sup>1</sup> Nevertheless, we chose to focus on the theoretical derivation of our method that and decided to limit it to separable metrics such as the isotropic  $L^2$  norm, in which case, (5) may be reformulated as

$$E(u^L, u^a, u^b) = \frac{1}{2} \sum_{c \in \{L, a, b\}} \int_{\Omega} (|\mathcal{D}^c| + \alpha |\nabla \mathcal{D}^c|) d\Omega \quad (6)$$

where  $\mathcal{D}^c$  is the  $c \in \{L, a, b\}$  color plane of  $\mathcal{D}$ .

Taking the first variation of (6) with respect to  $u^c$  for each  $c \in \{L, a, b\}$ , we get three EL equations (when discrete functions are involved, this process parallels the notion of zeroing the first derivative in order to minimize the function—see [3] for more details)

$$\begin{aligned} 0 &= \frac{\delta E(u)}{\delta u} = g * (-\alpha \nabla \cdot [\nabla g * (u^c - u_0^c)]) + g * (u^c - u_0^c) \\ &= g * (-\alpha \Delta + 1) * g * (u^c - u_0^c) \end{aligned} \quad (7)$$

(see Appendix A for this derivation). Reformulating the EL as a gradient descent flow for  $u^c$ , we get the following minimization scheme:

$$\left\{ \frac{du^c}{dt} = -g * (-\alpha \Delta + 1) \mathcal{D}^c \right\}_{c \in \{L, a, b\}} \quad \text{subject to } u^c \in \mathcal{G}. \quad (8)$$

These equations suggest that  $u^c$  should be iteratively updated in the reverse direction of the gradient of the penalty function defined in (6). This way, each iteration descends on the functional's surface, going to the minimum point of (6).

In the sequel, we will detail the conditions under which the optimum is unique. This would assure convergence to the single optimal solution. Note that (8) is a differential equation with an additional time parameter. The discretization of this parameter constitutes the index of iterations required to minimize (6) subject to  $u \in \mathcal{G}$  as a differential (8).

A direct consequence of the operators involved is the fact that the proposed functional and the resulting minimization scheme are both Euclidean invariant in the image plane. They are both translation and rotation invariant. As the parameter  $\alpha$  goes to zero, we approximate the S-CIELAB model, while for effective

<sup>1</sup>Alternatively, one may want to consider metrics that vary in color space.

nonzero and positive  $\alpha$ , we have a proper extension to the perceptual measures proposed in [14].

### B. Numerical Implementation

The above description concentrated on the model we use to force both gamut match and edge preservation. In this section, we shall describe the details of the numerical method to efficiently minimize (6).

In (8), we added an artificial time parameter  $t$  to the image  $u^c(x, y)$  that now reads  $u^c(x, y; t)$ . Let us discretize the EL gradient descent equation by first taking a simple forward explicit approximation for the  $t$  derivative. In the following, we shall refer to a single color layer ( $L$ ,  $a$ , or  $b$ ), as the same is to apply to all. Thus

$$\frac{u^{n+1} - u^n}{\tau} = \alpha g * \Delta \mathcal{D} - g * \mathcal{D} = g * (\alpha \Delta - 1) * g * (u - u_0)$$

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

Next, we deal with the spatial derivatives. Assume spatial discrete representation as well, where  $u_{i,j}^n \approx u(ih, jh; n\tau)$ , and assuming a uniform spatial spacings in the  $x$  and  $y$  directions of size  $h$ . We use central derivatives in space

$$u_{xx} \approx D_{xx}u \equiv \frac{u_{i+1,j} - 2u_{i,j} + u_{i-1,j}}{h^2}$$

$$u_x \approx D_x u \equiv \frac{u_{i+1,j} - u_{i-1,j}}{2h}$$

and a similar form for the  $y$  direction derivatives. 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

$$\begin{aligned} \tilde{\mathcal{D}}^n &= g * g * (u^n - u_0) \\ \mathcal{L}^n &= D_2 * (u^n - u_0) \\ u_{i,j}^{n+1} &= u_{i,j}^n + \tau \left( \alpha \mathcal{L}_{i,j}^n - \tilde{\mathcal{D}}_{i,j}^n \right) \end{aligned} \quad (9)$$

subject to the constraint  $u_{i,j}^n \in \mathcal{G}$ . Thus, every iteration constitutes of an update as in (9), followed by a projection onto the gamut to force the constraint.

In order to speed up convergence, we use a standard coarse to fine pyramidal approach. Thus, a Gaussian pyramid of the original image  $u_0$  is composed, and the gamut-mapping problem is solved for the coarse resolution. Since this stage works with very small images, the result is obtained very fast. Going to the next finer resolution layer, we solve again the gamut mapping problem, using an interpolated version of the previous resolution layer result as initialization. Again, due to the good initialization, convergence is rapid and the solution is obtained after only few (2–3) iterations. This process repeats until arrival to the finest resolution layer.

There are two ways to perceive this pyramidal treatment. One way is to think of it as a method that constructs a good candidate initialization to be used to solve the gamut mapping problem in the original (finest) resolution. Indeed, as we shall see next, if the gamut is convex and sufficient iterations are applied in each resolution layer, global optimum solution is guaranteed. Another, more instructive, way to interpret this pyramidal algorithm is as multiscale regularization. If we deliberately apply

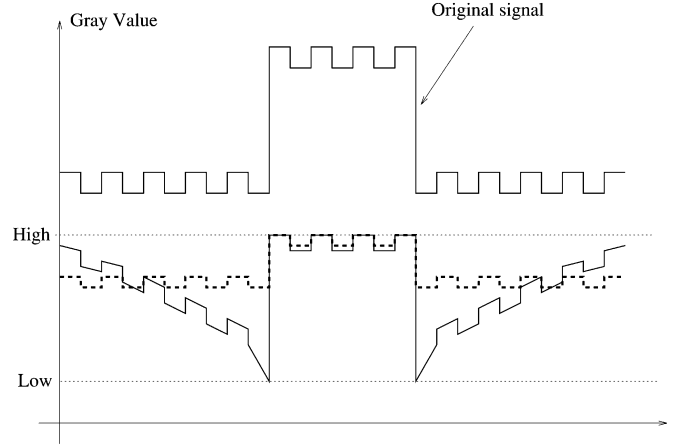


Fig. 2. One-dimensional example of the algorithm's behavior and the creation of halos.

only small number of iterations at each resolution layer, a desired smoothing effect is propagated through the pyramid. This has a lot in common with the the approach McCann advocated [13]. A third option exists, using a full multigrid method, but we have chosen not to pursue it in this sequel.

### C. 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 gamut set  $\mathcal{G}$  is convex, the overall problem is convex if and only if the Hessian of the functional is positive definite [3]. 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 and well posed for all  $\alpha > 0$ . Thus, for a convex target gamut  $\mathcal{G}$ , there exists a unique solution, and sufficient number of iterations will recover it. However, as was mentioned before, using a truncated number of iterations may be thought of as leading to a desired regularization effect, as is known in image restoration [11].

### D. Relation to Retinex

The gamut mapping problem is related to the Retinex problem of illumination compensation and dynamic range compression. The basic Retinex problem deals with the estimation of the reflectance image from the given acquired image. An optical model of the acquired image  $S$  asserts that it is a multiplication of the reflectance  $R$  and the illumination  $L$  images, where the reflectance image is a hypothetical 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, we require  $R < 1$ , implying  $r < 0$ . Thus, we seek an image  $r < 0$ , which is perceptually similar to  $s$ . For the Retinex problem 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



Fig. 3. Left to right: Original image, gamut mapping by truncation (minimization of the  $L_2$  norm), and the result of the proposed scheme.

just another way of saying 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 which also has perceptual features similar to  $u_0$ .

#### E. Robustifying the Algorithm

The proposed penalty function (5) tends to create halos in the resulting image. Fig. 2 explains the origin of those halos through a one-dimensional example. In this figure, we see a signal which is outside of the gamut (the gamut is between the “high” and “low” dotted lines, and the signal is strictly above the “high”). 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 halos, which means that, near the edge, there is a slow transition from low to high values.

In order to avoid this phenomena, we can modify the penalty term (5) using robust estimation formulations such as

$$E(u) = \int_{\Omega} \rho_1(\mathcal{D}) + \alpha \rho_2(|\nabla \mathcal{D}|) d\Omega \quad (10)$$

which, for  $\rho_1(x) = \rho_2(x) = x^2$ , coincides with (6). If the function  $\rho(x)$  grows slower than  $x^2$  as  $x \rightarrow \infty$ , we get an improved behavior near strong edges. Good candidates for  $\rho(x)$  are  $\rho(x) = |x|$  or  $\rho(x) = \sqrt{1 + x^2}$  (see [7] for more details on using these functions as replacements to the classic  $L^2$  norm). Note that by this change we have effectively changed the norm and, thus, the projection implied.

A different simpler (linear) approach with similar robust behavior is to solve the original problem (5) twice, with two different values of  $\alpha$ . We denote the solution with a small  $\alpha$  as  $u_{\text{small}}$  and the one which corresponds to the high value of  $\alpha$  as  $u_{\text{high}}$ . The solution  $u_{\text{small}}$  has smaller contrast at areas with



Fig. 4. (Top) Left to right: Original lighthouse image, gamut mapping by the original penalty function with  $\alpha = 1$ , and the same with  $\alpha = 40$ . (Bottom) Left to right: The weight image (white = 1, black = 0) and the final weighted average image.

small details, yet has almost no halos. On the other hand,  $u_{\text{high}}$  preserves the small details, but at the expense of strong halo effects. Therefore, by averaging these two results in a spatially adaptive way, we can enjoy both worlds. The proposed solution is, therefore

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

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.

Halo problems have been recently dealt with in relation to dynamic range compression. Solutions proposed included anisotropic diffusion [18], robust filtering [4], and bilateral filtering [5], [17]. See [2] and [7], [15] for a connection between these approaches and [10] and [16] for the relation to the Beltrami flow.

#### IV. RESULTS

The scope of the paper influenced our choice of experiments. We found that variations in addressing color fidelity and projection type are usually very dominant in determining the final image quality. Since our goal in this paper is the advocacy of a systematic spatial modeling via a variational point of view, we chose accordingly to restrict the underlying projection approach of the variational method to the simplest (orthogonal) projection. As a result, this is also the alternative method we compare

against. Thus, the results presented in this section are illustrative rather than extensive. Accordingly, we chose a toy problem with very limited gamut, rather than using a realistic one. A side benefit to this choice of gamut is our ability to magnify visual differentiation between the spatial and the nonspatial approaches and bypass possible limitations posed by the journal's production process.

We work in the RGB domain, and map into a toy gamut restricting the RGB values to the range [40, 100] out of the original gamut range [0, 255]. The projection norms we use in (6) are orthogonal in the RGB space are isotropic, and imply orthogonal projections. This is also the nonspatial projection we compare to.

Fig. 3 shows the result of the proposed mapping compared to nonspatial orthogonal projection. 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  $15 \times 15$  pixels.<sup>2</sup>

Next, we present the robust gamut mapping results. The applied algorithm is the shortcut method of adaptive weighting two regular results obtained with different values of  $\alpha$ . The top row of Fig. 4 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]. The bottom row in Fig. 4 shows the weight image as computed by the proposed formula with  $\beta = 0.005$ , and the weighted average result. As can be seen, the final result is a good tradeoff between halo suppression and sharpness.

<sup>2</sup>The size of the images is  $384 \times 256$  pixels.

## V. CONCLUDING REMARKS

We presented a variational formulation for the gamut mapping as a quadratic programming 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 of the difference between the images. We linked our results to previous methods including solutions to the Retinex problem, and presented an efficient numerical multiresolution algorithm for its solution, which can be used for image reproduction subject to convex constraints with a unique solution.

One important issue left untreated in this paper is the need to choose parameters. While being complex in general, we believe that the choice of the parameters should be partly driven by the desire to get scale-invariance. When the same image is fed to the algorithm in different scales, we could assume that the viewer expects the transition width near edges to be proportional to the image size. We leave this matter for future work.

## APPENDIX A

Let us explore the effect of a convolution operation within the functional on the EL equation. First, by linearity, it is simple to show that for a symmetric kernel  $g$  we have  $(d/dx)(g * u) = ((d/dx)g) * u = g * ((d/dx)u)$ , or, in shorthand notations,  $(d/dx)(g * u) = g_x * u$ . 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.,  $\int wv' = -\int w'v$ , 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  $\epsilon \rightarrow 0$ , such that  $dE/d\epsilon = 0$  for all  $\eta(x)$ . It is given by the EL equation

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

or, equivalently,  $g_x * F'(g_x * u) = 0$ . For example, for the functional

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

the EL is given by

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

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

## ACKNOWLEDGMENT

The authors would like to thank N. Moroney from Hewlett-Packard Laboratories, Palo Alto, CA, for pointing out [14], and A. Levi from Hewlett-Packard Laboratories, Haifa, Israel, for intriguing discussions on Sobolev Spaces and the Gamut problem during the summer of 1999.

## REFERENCES

- [1] R. Bala, R. deQueiroz, R. Eschbach, and W. Wu, "Gamut mapping to preserve spatial luminance variations," *J. Imag. Sci. Technol.*, vol. 45, pp. 436–482, 2001.
- [2] D. Barash, "Bilateral filtering and anisotropic diffusion: Toward a unified viewpoint," Hewlett-Packard Laboratories, Tech. Rep. 18-2000, Feb. 2000.
- [3] D. P. Bertsekas, *Non-Linear Programming*. Belmont, WA: Athena Scientific, 1995.
- [4] J. M. DiCarlo and B. A. Wandell, "Rendering high dynamic range images," presented at the Electronic Imaging Conf., Jan. 2000.
- [5] F. Durand and J. Dorsey, "Fast bilateral filtering for the display of high-dynamic range images," presented at the ACM SIGGRAPH Conf., Jul. 2002.
- [6] M. Elad, R. Kimmel, D. Shaked, and R. Keshet, "Reduced complexity retinex algorithm via the variational approach," Hewlett-Packard Laboratories, Tech. Rep. 19-2000, Feb. 2000.
- [7] M. Elad, "On the bilateral filters and ways to improve it," *IEEE Trans. Image Process.*, to be published.
- [8] R. Kimmel, M. Elad, D. Shaked, R. Keshet, and I. Sobel, "A variational framework for Retinex," *Int. J. Comput. Vis.*, to be published.
- [9] —, "A variational framework for Retinex," presented at the SPIE Electronic Imaging Conf., San Jose, CA, Jan. 2002.
- [10] R. Kimmel, R. Malladi, and N. Sochen, "Images as embedded maps and minimal surfaces: Movies, color, texture, and volumetric medical images," *Int. J. Comput. Vis.*, vol. 39, no. 2, pp. 111–129, Sep. 2000.
- [11] R. L. Lagendijk, *Iterative Identification and Restoration of Images*. New York: Kluwer, 1991.
- [12] S. Mallat, *A Wavelet Tour of Signal Processing*, 2nd ed. New York: Academic, 1998.
- [13] J. J. McCann, "Lessons learned from Mondrian applied to real images and color Gamuts," in *Proc. IS&T/SID 7th Color Imaging Conf.*, 1999, pp. 1–8.
- [14] S. Nakauchi, S. Hatanaka, and S. Usui, "Color Gamut mapping based on a perceptual image difference measure," *Color Res. Appl.*, vol. 24, pp. 280–291, 1999.
- [15] N. Sochen, R. Kimmel, and A. M. Bruckstein, "Diffusions and confusions in signal and image processing," *J. Math. Imag. Vis.*, vol. 14, no. 3, pp. 195–209, 2001.
- [16] N. Sochen, R. Kimmel, and R. Malladi, "A geometrical framework for low level vision," *IEEE Trans. Image Process.*, vol. 7, no. 3, pp. 310–318, Mar. 1998.
- [17] C. Tomasi and R. Manduchi, "Bilateral filtering for gray and color images," in *Int. Conf. Comput. Vis.*, Bombay, India, 1998, pp. 839–846.

- [18] J. Tumblin and G. Turk, "LCIS: A boundary hierarchy for detail-preserving contrast reduction," presented at the ACM SIGGRAPH Conf., 1999.
- [19] X. Zhang and B. A. Wandell, "A spatial extension of CIELAB for digital color image reproduction," in *Proc. SID Symp.*, vol. 27, 1996, pp. 731–734.



**Ron Kimmel** received the Ph.D. and D.Sc. degrees from The Technion—Israel Institute of Technology, Haifa, in 1996.

He is a Researcher in the areas of computer vision and image processing. He spent his postdoctoral years (from 1996 to 1998) at the University of California at Berkeley. He was a Visiting Professor at Stanford University, Stanford, CA, from 2003 to 2004, and a tenured Associate Professor at The Technion—Israel Institute of Technology. He has published over 100 articles and papers in scientific journals and conferences. He is on the editorial board of *International Journal of Computer Vision*, among others. He is the author of *Numerical Geometry of Images* (New York: Springer, 2003). His recent face recognition project (together with his students M. Bronstein and A. Bronstein) has been featured by CNN, WNBC, The Washington Post, Reuters, and more than 60 other newspapers and TV channels around the world. He was a Consultant to Hewlett-Packard Research Laboratories from 1998 to 2000 and Net2Wireless/Jigami research from 2000 to 2001, and he was on the advisory board of MediGuide (biomedical imaging, 2002 to 2004).

Prof. Kimmel is the recipient of the Rich Technion Innovation Award (twice), the Taub Prize for Excellence in Research, and the Alon Fellowship. He is on the editorial board of the IEEE TRANSACTIONS ON IMAGE PROCESSING.



**Doron Shaked** graduated from the Electrical and Computer Engineering Department, Ben Gurion University, Beer Sheva, Israel, in 1988, and received the M.Sc. and D.Sc. degrees from the Electrical Engineering Department, The Technion—Israel Institute of Technology, Haifa, in 1991 and 1995, respectively.

He is a Senior Researcher at Hewlett Packard Laboratories, Haifa, Israel, where he has been since 1995, working on image processing and analysis, with a special interest in printing technologies.

Dr. Shaked was an Ollendorf Student Fellow in 1991 and the recipient of the Wolf prize for excellent students in 1994.



**Michael Elad** received the B.Sc., M.Sc., and D.Sc. degrees from the Department of Electrical Engineering at The Technion—Israel Institute of Technology (IIT), Haifa, Israel, in 1986, 1988, and 1997, respectively.

From 1988 to 1993, he served in the Israeli Air Force. From 1997 to 2000, he worked at Hewlett-Packard Laboratories as an R&D Engineer. From 2000 to 2001, he headed the research division at Jigami Corporation, Israel. From 2001 to 2003, he was a Research Associate with the Computer Science Department, Stanford University (SCCM program), Stanford, CA. In September 2003, he joined the Department of Computer Science, IIT, as an Assistant Professor. He was also a Research Associate at IIT from 1998 to 2000, teaching courses in the Electrical Engineering Department. He works in the field of signal and image processing, specializing, in particular, on inverse problems, sparse representations, and over-complete transforms.

Dr. Elad received the Best Lecturer Award twice (in 1999 and 2000). He is also the recipient of the Guttwirth and the Wolf fellowships.



**Irwin Sobel** received the Ph.D. degree in electrical engineering and computer science from Stanford University, Stanford, CA, and the B.S. degree in mathematics and electrical engineering from the Massachusetts Institute of Technology, Cambridge.

He has been a Senior Researcher at Hewlett-Packard Laboratories, Palo Alto, CA, since December 1982, where he has worked on color image processing, computer vision, and three-dimensional visualization. Prior to this, he spent nine years at the Department of Biological Sciences, Columbia

University, New York, automating serial-section reconstruction for neuro-anatomy—the precursor of this project.