색상 양자화된 영상의 복원

김태훈* 최민규† 안종우‡°
*삼성전자 †‡광운대학교
*thn.kim@samsung.com †mgchoi@kw.ac.kr ‡jahn@cs.kw.ac.kr

Restoration of Color-Quantized Images

Tae-hoon Kim* Min Gyu Choi† Jongwoo Ahn‡°
*Samsung Electronics Co., Ltd. †‡Kwangwoon University

Abstract

Color quantization replaces the color of each pixel with the closest representative color, and thus it makes the resulting image partitioned into uniformly-colored regions. As a consequence, continuous, detailed variations of color over the corresponding regions in the original image are lost through color quantization. In this paper, we present a novel scheme for restoring such variations from a color-quantized input image. Our scheme identifies which pairs of uniformly-colored regions in the input image should have continuous variations of color in the resulting image. Then, such regions are seamlessly stitched using the Laplace equation. The user can optionally indicate which regions should be separated or stitched by scribbling *constraint brushes* across the regions. We demonstrate the effectiveness of our approach through diverse examples, such as photographs, cartoons, and artistic illustrations.

1 Introduction

Color quantization is a lossy process that reduces the number of colors in an image with a minimal visual artifact. Thus a color-quantized image can be regarded as a degraded version of its original. There have been many researches to restore or enhance visual qualities of given images such as deblurring [2], noise reduction [8, 13], spatial super-resolution [6, 11], image completion [3, 5], and colorization [9, 14]. However, there has been a little effort to restore a color-quantized image. We call this restoration *image dequantization*. Image dequantization is important because it can increase the visual quality of an image comprised of a small number of colors, such as images acquired from devices with limited color capabilities or transmitted via narrow-band networks (See Figure 1).

According to the color quantization theory, each color in the original image is quantized to its closest representative color. This quantization makes the resulting image partitioned into a set of non-overlapping, uniformly-colored regions, as illustrated in Figure 2. Thus detailed variations of color in the regions as well as across the neighboring regions are lost through color quantization. Then, image dequantization, the inverse operation of color quantization, calls for restoration of such continuous variations of color. In this paper, we present a novel scheme for image dequantization. Our scheme first identifies which pairs of uniformly-colored regions in the input image should have continuous variations in the resulting image. Then, such regions are seamlessly stitched based on the Laplace equation. The user can also indicate which regions should be separated or stitched seamlessly by scribbling constraint brushes across the regions.

1.1 Related Work

The problem of reducing degradation and noise in images has been addressed for a long time. Early approaches have used spatial filtering techniques such as median and Winer filter. These approaches were extended to Kalman filter for multiple channels [8] and a fuzzy smoothing operation [12]. Beyond reducing a noise channel, researches have been focused on reconstructing images of higher spatial resolution based on the observed features in the input images. Borman and Stevenson [4] have combined multiple low-resolution images obtained at sub-pixel displacements. Freeman et al. [6] have exploited the stored high-resolution patch corresponding to every possible low-resolution image patch. Baker and Kanade [1] have enhanced salient features recognized in the low resolution images.

Another important issue is completing irregular missing portions caused by removing foreground or background objects from an image. Bertalmio et al. [3] have used PDE-based approach and Drori et al. [5] have adopted texture synthesis at the image patch level. Colorizing a gray-scale image has also been a topic of considerable interest. Welsh et al. [14] have determined color

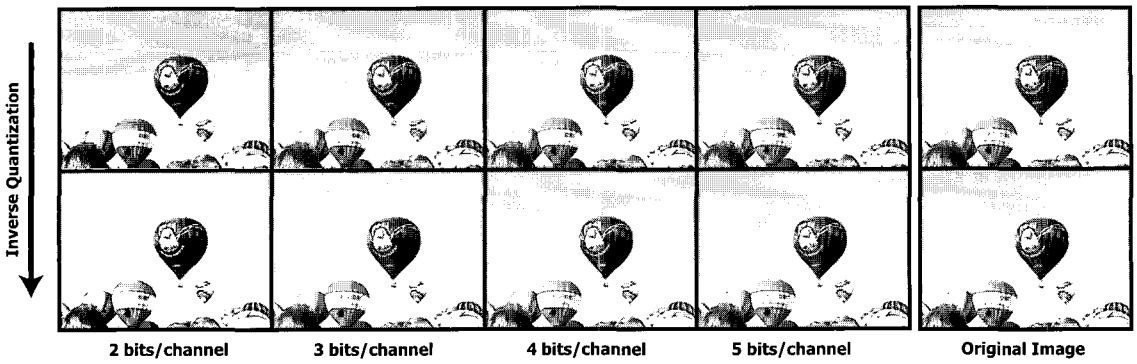

Figure 1: Image dequantization restores continuous variations of color in a quantized image. The images (top row) with a small number of bits per channel are restored using our method so as to exhibit continuous variations of color in the resulting images (bottom row).

of each pixel from those with matching neighborhoods in examples. Levin et al. [9] have formulated an optimization problem based on the premise that neighboring pixels with similar intensities should have similar colors.

Although a lot of approaches for image restoration have been proposed, to our knowledge, there have been a few attempts to automatically recover an original image from a given color-quantized image. For image quantization with two colors, that is, *dithering*, noise reduction methods have been successively applied [13]. For more than two colors, Fung and Chan [7] developed a regularized method that iteratively refines a given quantized image so that each pixel has a similar color with its neighboring pixels while preserving its closest representative color. However, this method attempts to seamlessly stitch adjacent regions that should be separated, because smoothing for a pixel involves all of its neighboring pixels that may belong to different objects.

2 Inverse Quantization

2.1 Observation

In color quantization of an image, a set of its representative color vectors is firstly selected. Then, the color vector of each pixel in the image is quantized to its closest one from the set of representative color vectors, based on the Euclidean distance. Consequently, the color space is partitioned into a set of non-overlapping Voronoi cells, each of which corresponds to the representative color. This color quantization also partitions the resulting image into a set of non-overlapping regions, each of which consists of pixels of the same color, as illustrated in Figure 2. Thus, the detailed variation of color in each region is lost through color quantization. Then, image dequantization, the inverse operation of color quantization, calls for restoration of such a continuous variation of color in each uniformly-colored region of the quantized image.

Before addressing dequantization of uniformly-colored regions, we first need to examine their adjacency relationships not only in the image space but also in the color space. For a pair of uniformly-colored regions adjacent in the image space, their corresponding Voronoi cells in the color space can be either (1) adjacent or (2) not adjacent. In the first case, we would expect a seamless variation of color across the edge between the two regions. However, in the second case, discontinuities at the edge would be preferred because it could be thought of as an edge in the original image. We call the first type of edge *soft* and the second *hard*. Note that soft edges are caused by quantization whereas hard edges have been existing before quantization.

We now examine pixels on soft edges more precisely. Suppose that two neighboring pixels induces a soft edge in the image space as illustrated in Figure 3. Then, their original colors must have been located near the Voronoi edge between their corresponding Voronoi cells in the color space. Thus, it is natural to infer the unknown original colors from the average of their quantized colors. In contrast, for pixels on hard edges, we assume that their original colors and quantized colors would be the same.

Based on the above observations, we conclude that image dequantization is to enforce continuous variations of color not only over uniformly-colored regions but also across soft edges, while preserving discontinuities at hard edges. For continuous variations of color, we are to employ Laplace's equation with Dirichlet boundary conditions. Recall that pixels inducing a hard edge are desired to retain their quantized colors, and pixels inducing a soft edge are desired to have the average of their quantized colors. These will be used as Dirichlet boundary conditions.

2.2 Formulation with Laplace's Equation

Before introducing image dequantization using Laplace's equation, we first define necessary notations. For a pixel $\mathbf{p}$, let $N(\mathbf{p})$ be a set of its 4-connected neighbors. Then, the boundary ∂R_i of a uniformly-colored region R_i can be defined as follows: For

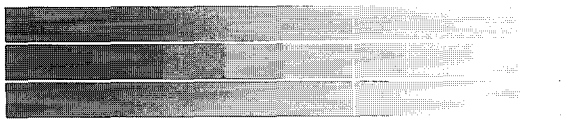

Figure 2: An image with red to green gradation (top) is quantized with 8 colors (middle). Our image dequantization algorithm (bottom) restores the original gradation successfully.

a pair of pixels $(\mathbf{p}, \mathbf{q})$ such that $\mathbf{p} \in R_i$, $\mathbf{q} \in N(\mathbf{p})$ but $\mathbf{q} \notin R_i$, we introduce a virtual pixel $\mathbf{v}_{\mathbf{p},\mathbf{q}} = (\mathbf{p} + \mathbf{q})/2$ and let it belong to the boundary ∂R_i. If the Voronoi cells corresponding to the pixels $\mathbf{p}$ and $\mathbf{q}$ are adjacent in the color space, then the virtual pixel $\mathbf{v}_{\mathbf{p},\mathbf{q}}$ is interpreted as a part of a soft edge (triangular pixels in Figure 3). Otherwise, it is interpreted as a part of a hard edge (square pixels in Figure 3). Thus, the boundary ∂R_i consist of soft edges and hard edges, and it can be thought of as an interface between neighboring regions.

Now, we detail image dequantization using Laplace's equation. As it is enough to solve the image dequantization problem for each color channel independently, we consider only scalar image functions. For a uniformly-colored region R_i, let φ be the unknown scalar function to be defined over $R_i \cup \partial R_i$. We first define a scalar function φ^* over ∂R_i for boundary conditions: if a virtual pixel $\mathbf{v}_{\mathbf{p},\mathbf{q}}$ is on a hard edge of ∂R_i, $\varphi^*(\mathbf{v}_{\mathbf{p},\mathbf{q}})$ is set to the current, quantized color of $\mathbf{p} \in R_i$; otherwise, $\varphi^*(\mathbf{v}_{\mathbf{p},\mathbf{q}})$ is set to the average color of the pixels $\mathbf{p}$ and $\mathbf{q}$. This averaging allows continuous variations of color across the three consecutive pixels $\mathbf{p}$, $\mathbf{v}_{\mathbf{p},\mathbf{q}}$, and $\mathbf{q}$. Finally, we fill the region R_i with a continuous variation of color by solving the following Euler-Lagrange equation:

$$\Delta\varphi = 0 \text{ over } R_i \text{ with } \varphi|_{\partial R_i} = \varphi^*|_{\partial R_i}, \tag{1}$$

where $\Delta\cdot = \frac{\partial^2 \cdot}{\partial x^2} + \frac{\partial^2 \cdot}{\partial y^2}$ is the Laplacian operator.[1]

To build a linear system for Equation (1), we need to compute the second partial derivative of φ with respect to x and also with respect to y. A virtual pixel on the boundary ∂R_i is 0.5 pixel apart from a pixel in the region R_i along either the x- or y-direction in the image space. Suppose that three pixels $\mathbf{p}$, $\mathbf{q}$, $\mathbf{r}$ are consecutive along the x-direction. When only $\mathbf{r}$ is a virtual pixel on the boundary ∂R_i for instance, the step sizes for differentiation are 1.0 between $\mathbf{p}$ and $\mathbf{q}$ and 0.5 between $\mathbf{q}$ and $\mathbf{r}$. The derivative at $\mathbf{q}$ can be obtained by differentiating the Lagrange interpolation polynomial that passes through $(-1.0, \varphi(\mathbf{p}))$, $(0, \varphi(\mathbf{q}))$, and $(0.5, \varphi(\mathbf{r}))$. Summarizing all the cases, the second partial

[1] Three Laplace equations of the form (1) are solved independently in the three color channels of the chosen color space. We have obtained all the results in the RGB color space because commonly available color-quantized images have been processed in that space. According to our experiments, similar results were obtained in the YUV color space for instance.

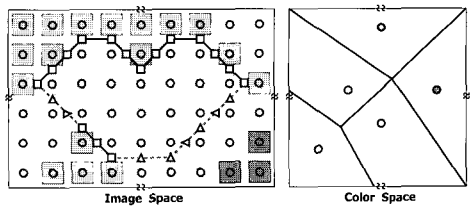

Figure 3: The boundary of the gray-colored region consists of soft edges (drawn in dashed lines) and hard edges (drawn in solid lines).

derivative is given as follows:

$$\left.\frac{\partial^2\varphi}{\partial x^2}\right|_{\mathbf{q}} = \begin{cases} \varphi(\mathbf{p}) - 2\varphi(\mathbf{q}) + \varphi(\mathbf{r}) & \text{if } \mathbf{p} \notin \partial R_i, \mathbf{r} \notin \partial R_i, \\ \frac{4}{3}\varphi(\mathbf{p}) - 4\varphi(\mathbf{q}) + \frac{8}{3}\varphi(\mathbf{r}) & \text{if } \mathbf{p} \notin \partial R_i, \mathbf{r} \in \partial R_i, \\ \frac{8}{3}\varphi(\mathbf{p}) - 4\varphi(\mathbf{q}) + \frac{4}{3}\varphi(\mathbf{r}) & \text{if } \mathbf{p} \in \partial R_i, \mathbf{r} \notin \partial R_i, \\ 4\varphi(\mathbf{p}) - 8\varphi(\mathbf{q}) + 4\varphi(\mathbf{r}) & \text{if } \mathbf{p} \in \partial R_i, \mathbf{r} \in \partial R_i. \end{cases} \tag{2}$$

The second partial derivative of φ with respect to y is of the same form except that three pixels are consecutive along the y-direction.

Applying the above formula to Equation (1) for the region R_i yields a linear system of the following form:

$$\mathbf{A}_i\mathbf{x}_i = \mathbf{b}_i, \tag{3}$$

where $\mathbf{x}_i$ consists of all the unknown image function values φ of the region R_i, and the matrix $\mathbf{A}_i$ is sparse, symmetric, and positive-definite. Here, we note that the image function values φ on the boundary ∂R_i are the same with φ^* as specified in Equation (1) and they are involved with only $\mathbf{b}_i$. As a consequence, we build the linear systems of the form (3) for all the regions independently, and then assemble them into a large, sparse, symmetric, and positive definite linear system so as to solve the system efficiently with standard methods employed in [9, 10].

An example of image dequantization is illustrated in Figure 2. We obtained a red to green gradation image using Adobe Photoshop and then quantized it with 8 colors. Our image dequantization method restores the original gradation successfully in that we can hardly notice the difference between the dequantized image and the original. In this example, there are only soft edges that allow continuous variations of colors. Figure 4 shows more complex examples, where the color-quantized input images have not only soft edges but also hard edges. We can observe that the dequantized images preserve discontinuities at hard edges successfully. Our method can also be used for network-based imaging with progressive n-bit quantization. In this scheme, consecutive n bits starting from the most significant one for each color channel are used for progressive improvement of the image. Figure 1 exhibits that our method produces an image with quite good quality even with 3 bits per channel for each pixel.

Figure 4: Comparison of color-quantized input, dequantized, and original images.

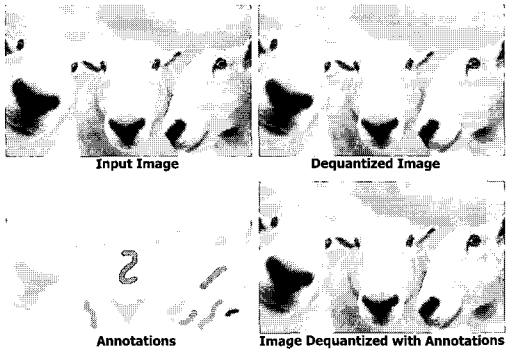

Figure 5: Interactive dequantization with annotations.

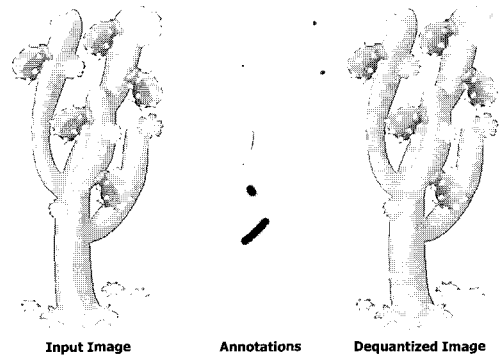

Figure 6: A cartoon-shaded image is dequantized interactively.

2.3 Interactive Annotation

The type of the edge between two neighboring regions has been classified into either soft or hard, solely based on the adjacency of their corresponding Voronoi cells in the color space. However, there are situations where it is desirable to override some adjacency relationships interactively. For example, in Figure 5, the fur of the sheep has very fine details even with quantized colors, however such fine details are lost in the dequantized image. This is because the Voronoi cells corresponding to the quantized colors of the fur are adjacent in the color space. We provide interactive annotation that overrides the adjacency relationships in the color space. The magenta brush is to annotate that the Voronoi cells corresponding to the quantized colors under the brush are not adjacent. The blue brush is to annotate the reverse; distant Voronoi cells are treated as if they are adjacent.

Interactive annotation can also be used for cartoons and artistic illustrations, in which artists carefully select representative colors to shade objects with similar but somewhat different colors. Figure 6 shows a cartoon image dequantized with annotations. In this example, we introduced new brushes that override the type of the edge only in the image space; the type of the edge annotated with a green (red) brush is turned into a hard (soft) edge. In cartoon-shaded images, hard edges can also be utilized for decoration of the resulting images, as illustrated in Figure 7.

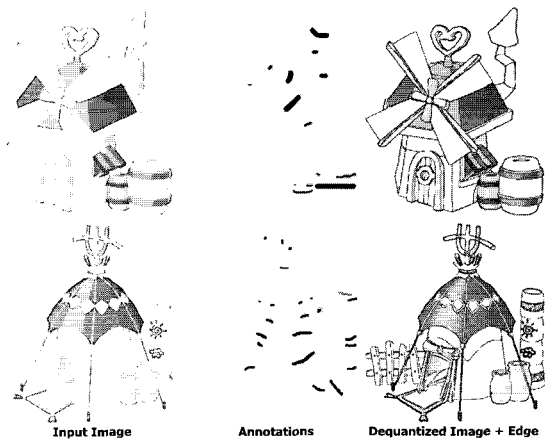

Figure 7: Hard edges are drawn for cartoon shading.

2.4 Continuity Constraints

Our image dequantization method guarantees only C^0 continuity between two neighboring regions. This limited continuity is not perceptible when the intensity change between the two regions is relatively small. However, when the intensity change is relatively large as the inside of the teapot in Figure 8, Laplace's equation with Dirichlet boundary conditions may produce an unsatisfactory result. The Mach band effect is observed inside the dequantized teapot. To reduce such an illusion, we introduce additional continuity constraints.

Suppose that $\mathbf{p}, \mathbf{v}_{\mathbf{p},\mathbf{q}}, \mathbf{q}$ are three pixels consecutive along the x-direction in the image space. As the notation indicates $\mathbf{v}_{pq}$ is the virtual pixel considered as a part of a soft edge. Then, we impose a new continuity constraint on the image function values at $\mathbf{p}$ and $\mathbf{q}$ such that the backward and forward differences of intensity at $\mathbf{v}_{pq}$ should be the same:

$$\varphi(\mathbf{p}) + \varphi(\mathbf{q}) = 2\varphi^*(\mathbf{v}_{\mathbf{p},\mathbf{q}}). \tag{4}$$

In the same vein, we also add continuity constraints along the y-direction. These additional constraints result in an overdetermined system of linear equations. By multiplying a weighting factor w to the constraint equations, we control the significance of continuity at boundaries, as illustrated in Figure 8.

3 Summary

We have presented a novel scheme for restoring a color-quantized input image that consists of uniformly-colored regions. Considering adjacency relationships of the regions in the image space as well as in the color space, our scheme classifies the region

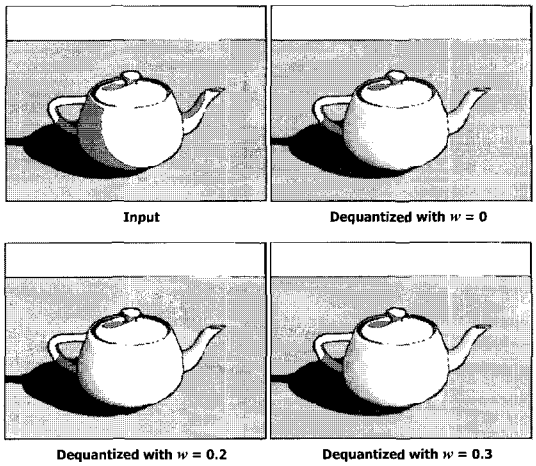

Figure 8: The Mach band illusion ($w = 0$) disappears through the use of additional continuity constraints ($w \neq 0$).

boundaries into soft and hard. Together with this classification, Laplace's equation is employed for continuous variations of color not only over uniformly-colored regions but also across soft boundaries between neighboring regions, while preserving discontinuities at hard boundaries. The user can optionally overrides the classification result by scribbling brushes across the regions. We have demonstrated the effectiveness of our approach through diverse examples, such as photographs, cartoons, and artistic illustrations.

Acknowledgment

This work was supported in part by Seoul Research and Business Development Program (10581) and Ministry of Information and Communication under the Information Technology Research Center (ITRC) support program.

References

[1] Simon Baker and Takeo Kanade. Hallucinating faces. In *Proc. the Fourth International Conference on Automatic Face and Gesture Recognition*, Grenoble, France, 2000.

[2] Moshe Ben-Ezra and Shree K. Nayar. Motion-based motion deblurring. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(6):689–698, 2004.

[3] Marcelo Bertalmio, Guillermo Sapiro, Vicent Caselles, and Coloma Ballester. Image inpainting. In *Proc. ACM SIGGRAPH*, pages 417–424, 2000.

[4] Sean Borman and Robert Stevenson. Spatial resolution enhancement of low-resolution image sequences - a comprehensive review with directions for future research. Technial report, Laboratory for Image and Signal Analysis (LISA), University of Notre Dame, 1998.

[5] Iddo Drori, Daniel Cohen-Or, and Hezy Yeshurun. Fragment-based image completion. *ACM Transactions on Graphics*, 22(3):303–312, 2003.

[6] Willam T. Freeman, Thouis R. Jones, and Egon C Pasztor. Example-based super-resolution. *IEEE Computer Graphics and Applications*, 22(2):56–65, 2002.

[7] Yik-Hing Fung and Yuk-Hee Chan. An iterative algorithm for restoring color-quantized images. In *Proc. International Conference on Image Processing*, pages 313–316, 2002.

[8] Nikolas P. Galatsanos and Roland T. Chin. Restoration of color images by multichannel Kalman filtering. *IEEE Transactions on Signal Processing*, 39(10):2237–2252, 1991.

[9] Anat Levin, Dani Lischinski, and Yair Weiss. Colorization using optimization. *ACM Transactions on Graphics*, 23(3):689–694, 2004.

[10] Patric Pérez, Michel Gangnet, and Andrew Blake. Poisson image editing. *ACM Transactions on Graphics*, 22(3):313–318, 2003.

[11] Eli Shechtman, Yaron Caspi, and Michael Irani. Space-time super-resolution. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(4):531–545, 2005.

[12] Dimitri Van De Ville, Mike Nachtegael, Dietrich Van der Weken, Etienne E. Kerre, Wilfried Philips, and Ignance Lemahieu. Noise reduction by fuzzy image filtering. *IEEE Transactions on Fuzzy Systems*, 11(4):429–436, 2003.

[13] Tsachy Weissman, Erik Ordentlich, and Gradiel Seroussi. Universal discrete denoising: Known channel. *IEEE Transactions on Information Theory*, 51(1):5–28, 2005.

[14] Tomihisa Welsh, Michael Ashikhmin, and Klaus Mueller. Transferring color to greyscale images. *ACM Transactions on Graphics*, 21(3):277–280, 2002.