

Image Dequantization: Restoration of Quantized Colors

Tae-hoon Kim^{†1}, Jongwoo Ahn², and Min Gyu Choi²

¹Olaworks, Inc., Korea

²Kwangwoon University, Korea

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 blind scheme for restoring such variations from a color-quantized input image without a priori knowledge of the quantization method. 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 through optimization while preserving the closest representative colors. 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.

Categories and Subject Descriptors (according to ACM CCS): I.4.3 [Image Processing and Computer Vision]: Enhancement – Filtering; I.4.9 [Image Processing and Computer Vision]: Applications

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 [BEN04], noise reduction [GC91, WOS05], spatial super-resolution [FJP02, SCI05], image completion [BSCB00, DCOY03], colorization [LLW04, WAM02], and reconstruction of a high dynamic range (HDR) image [LSA05]. 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 that had been acquired in the early days from devices with limited color capabilities or transmitted via narrow-band networks.

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 re-

gions, as illustrated in Figure 3. 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 *blind* image dequantization that does not require a priori knowledge on the quantization method of the input image. Our scheme first identifies which pairs of uniformly-colored regions in the input image should have continuous variations in the resulting image. Then, seamless stitching of such regions is formulated as an optimization problem. 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 [GC91] and a fuzzy smoothing operation [VNdW*03]. Beyond reducing a noise channel, researches have been focused on reconstructing images of

[†] thkim@olaworks.com, jwahn@cs.kw.ac.kr, mgchoi@kw.ac.kr

higher spatial resolution based on the observed features in the input images. Borman and Stevenson [BS98] have combined multiple low-resolution images obtained at sub-pixel displacements. Freeman et al. [FJP02] have exploited the stored high-resolution patch corresponding to every possible low-resolution image patch. Baker and Kanade [BK00] 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. [BSCB00] have used PDE-based approach and Drori et al. [DCOY03] 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. [WAM02] have determined color of each pixel from those with matching neighborhoods in examples. Levin et al. [LLW04] 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 [WOS05]. For more than two colors, Fung and Chan [FC02] 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. Li et al. [LSA05] presented the companding scheme that encodes an original HDR image into an LDR image that can be decoded as close as possible to the original image. Our work is distinguished from their approach in that, whereas the companding method can be applied only to the image encoded by the method itself, our method can be successfully applied to an image quantized with an unknown algorithm; our image dequantization method is universal.

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 1. Thus, the detailed vari-

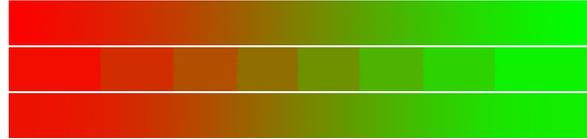


Figure 1: An image with red to green gradation (top) is quantized with 8 colors (middle). Our image dequantization algorithm (bottom) restores the original gradation successfully.

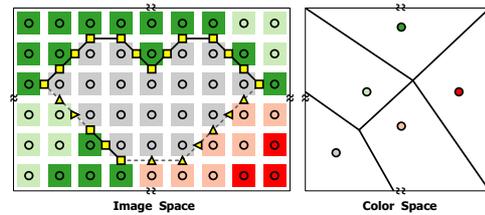


Figure 2: The boundary of the gray-colored region consists of soft edges (drawn in dashed lines) and hard edges (drawn in solid lines).

ation 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*. Then the image dequantization problem can be simplified by taking an assumption 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 induce a soft edge in the image space as illustrated in Figure 2. 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

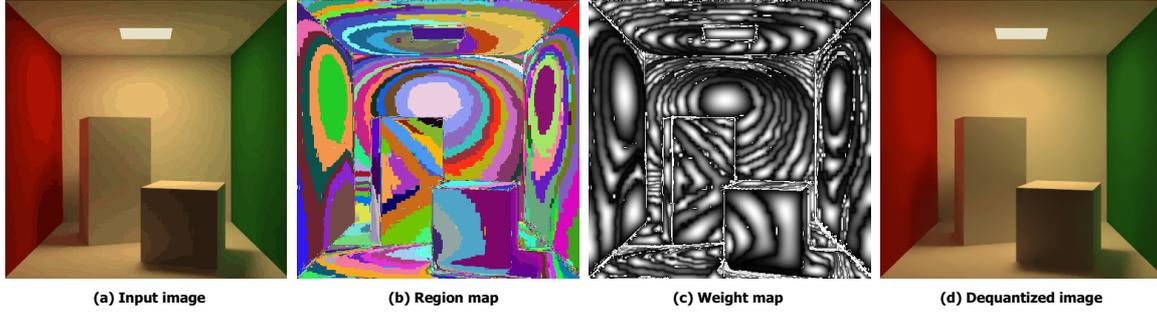


Figure 3: An example of image dequantization by optimization. (a) The input image. (b) An indexed image of uniformly-colored regions. (c) The weight function $w(\mathbf{p})$ based on the minimum distance to the boundaries (higher intensity implies a large weight). (d) The dequantized image obtained by optimization with the weight function.

not only over uniformly-colored regions but also across soft edges, while preserving discontinuities at hard edges and making pixels preserve their representative colors. For continuous variations of color and preservation of the representative color, we are to minimize an objective function containing a smoothness term and a data term. 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 boundary conditions.

2.2. Formulation

Prior to defining an objective function to minimize, we first fix necessary notations. Let R_i be a uniformly-colored region in the image space. Then the boundary ∂R_i of the region R_i is defined as follows: For a pixel \mathbf{p} and its 4-connected neighbor $\mathbf{q} \notin R_i$, 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 2). Otherwise, it is interpreted as a part of a hard edge (square pixels in Figure 2). Thus, the boundary ∂R_i consists of soft edges and hard edges, and it can be thought of as an interface between adjacent regions. For a pixel \mathbf{p} , we define $N_{\mathbf{p}}$ to be the set of its neighbors $\mathbf{q} \in R_i \cup \partial R_i$.

Now, we formulate image dequantization as an optimization problem. 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 $R_i \cup \partial R_i$ and $\bar{\varphi}$ be the scalar function corresponding to the input image, that is, $\bar{\varphi}(\mathbf{p})$ is the quantized color of \mathbf{p} . We first define a scalar function φ^* over ∂R_i for boundary

[†] We have obtained all the results in the CIE-Lab color space designed to approximate human visual system.

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 of $\bar{\varphi}(\mathbf{p})$ and $\bar{\varphi}(\mathbf{q})$. 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 by minimizing the following objective function:

$$J(\varphi) = \sum_{(\mathbf{p},\mathbf{q}) \cap R_i \neq \emptyset} (\varphi(\mathbf{p}) - \varphi(\mathbf{q}))^2 + \sum_{\mathbf{p} \in R_i} w(\mathbf{p}) (\varphi(\mathbf{p}) - \bar{\varphi}(\mathbf{p}))^2, \quad (1)$$

with $\varphi(\mathbf{v}_{\mathbf{p},\mathbf{q}}) = \varphi^*(\mathbf{v}_{\mathbf{p},\mathbf{q}})$ for every virtual pixel $\mathbf{v}_{\mathbf{p},\mathbf{q}} \in \partial R_i$. Here, (\mathbf{p}, \mathbf{q}) denotes a pair of pixels such that $\mathbf{q} \in N_{\mathbf{p}}$, and $w(\mathbf{p})$ is a weight function. The former term controls the smoothness of the pixels in R_i and the latter term is the data term that makes the pixels preserve their original colors. The smoothness term is more important at the pixels near the boundary ∂R_i while the data term is more important at the pixels far from the ∂R_i .

The weight function $w(\mathbf{p})$ is determined based on the minimum distance $d(\mathbf{p})$ from the pixel $\mathbf{p} \in R_i$ to the boundary ∂R_i , which can be computed by solving a Poisson equation of the form [GGG*04]:

$$\Delta d(\mathbf{p}) = 1 \text{ over } R_i \text{ with } d|_{\partial R_i} = 0, \quad (2)$$

where $\Delta \cdot = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}$ is the Laplacian operator. Then, $d(\mathbf{p})$ is normalized by applying a bell-shaped function:

$$w(\mathbf{p}) = \omega \cdot \left(\frac{1 - \exp(d(\mathbf{p})/d_{max})}{1 - \exp(1)} \right), \quad (3)$$

where $d_{max} = \max_{\mathbf{p} \in R_i} d(\mathbf{p})$, and ω is a constant controlling the global influence of the data term to $J(\varphi)$ in R_i . An example of $w(\mathbf{p})$ is shown in Figure 3(c).

The Equation (1) can be minimized by solving the following simultaneous linear equations:

$$|N_{\mathbf{p}}| \varphi(\mathbf{p}) - \sum_{\mathbf{q} \in N_{\mathbf{p}} \cap R_i} \varphi(\mathbf{q}) + w(\mathbf{p}) (\varphi(\mathbf{p}) - \bar{\varphi}(\mathbf{p})) = \sum_{\mathbf{q} \in N_{\mathbf{p}} \cap \partial R_i} \varphi^*(\mathbf{v}_{\mathbf{p},\mathbf{q}}),$$

for all $\mathbf{p} \in R_i$. Assembling the above equations for the region

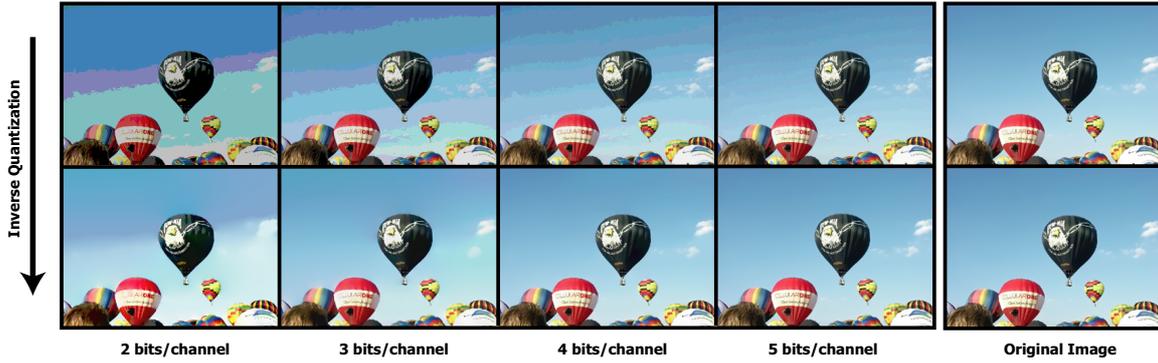


Figure 4: Image dequantization restores continuous variations of color in a quantized image. The images (top row) with a small number of bits per channel are restored using our method so as to exhibit continuous variations of color in the resulting images (bottom row).

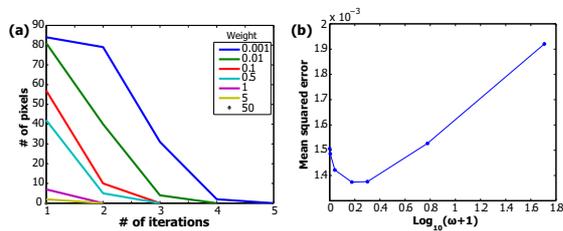


Figure 5: Dequantization with various weights. (a) The number of pixels out of their Voronoi cells vs. the number of iterations. (b) Mean squared error vs. log-scaled weight.



Figure 6: Weight vs. visual quality. (a) The image with the minimum MSE. (b) The image with the best visual quality.

R_i yields a linear system of the following form:

$$\mathbf{A}_i \mathbf{x}_i = \mathbf{b}_i, \quad (4)$$

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 can build the linear systems of the form (4) 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 [LLW04, PGB03].

An example of image dequantization is illustrated in Figure 1. 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 color. Figure 3 shows a more complex example of our scheme, 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 4 exhibits that our method produces an image with quite good quality even with 3 bits per channel for each pixel.

2.3. Iterative Refining

The dequantized image φ obtained by solving Equation (4) may not be invertible, that is, re-quantization of φ with the representative colors of the input image $\bar{\varphi}$ may be different from φ , since we can not guarantee $\varphi(\mathbf{p})$ to be always in the same Voronoi cell corresponding to $\bar{\varphi}(\mathbf{p})$.

We can simply decrease the likelihood of \mathbf{p} to be out of its Voronoi cell by increasing the weight function $w(\mathbf{p})$ in Equation (1). This strategy can be used to iteratively refine the dequantized image: We first compute φ and then re-quantize φ to φ^R . For each pixel \mathbf{p} out of its Voronoi cell, that is, $\varphi^R(\mathbf{p}) \neq \bar{\varphi}(\mathbf{p})$, we increase $w(\mathbf{p})$. We repeat these steps until φ^R is identical to $\bar{\varphi}$. Although the weight function $w(\mathbf{p})$ may

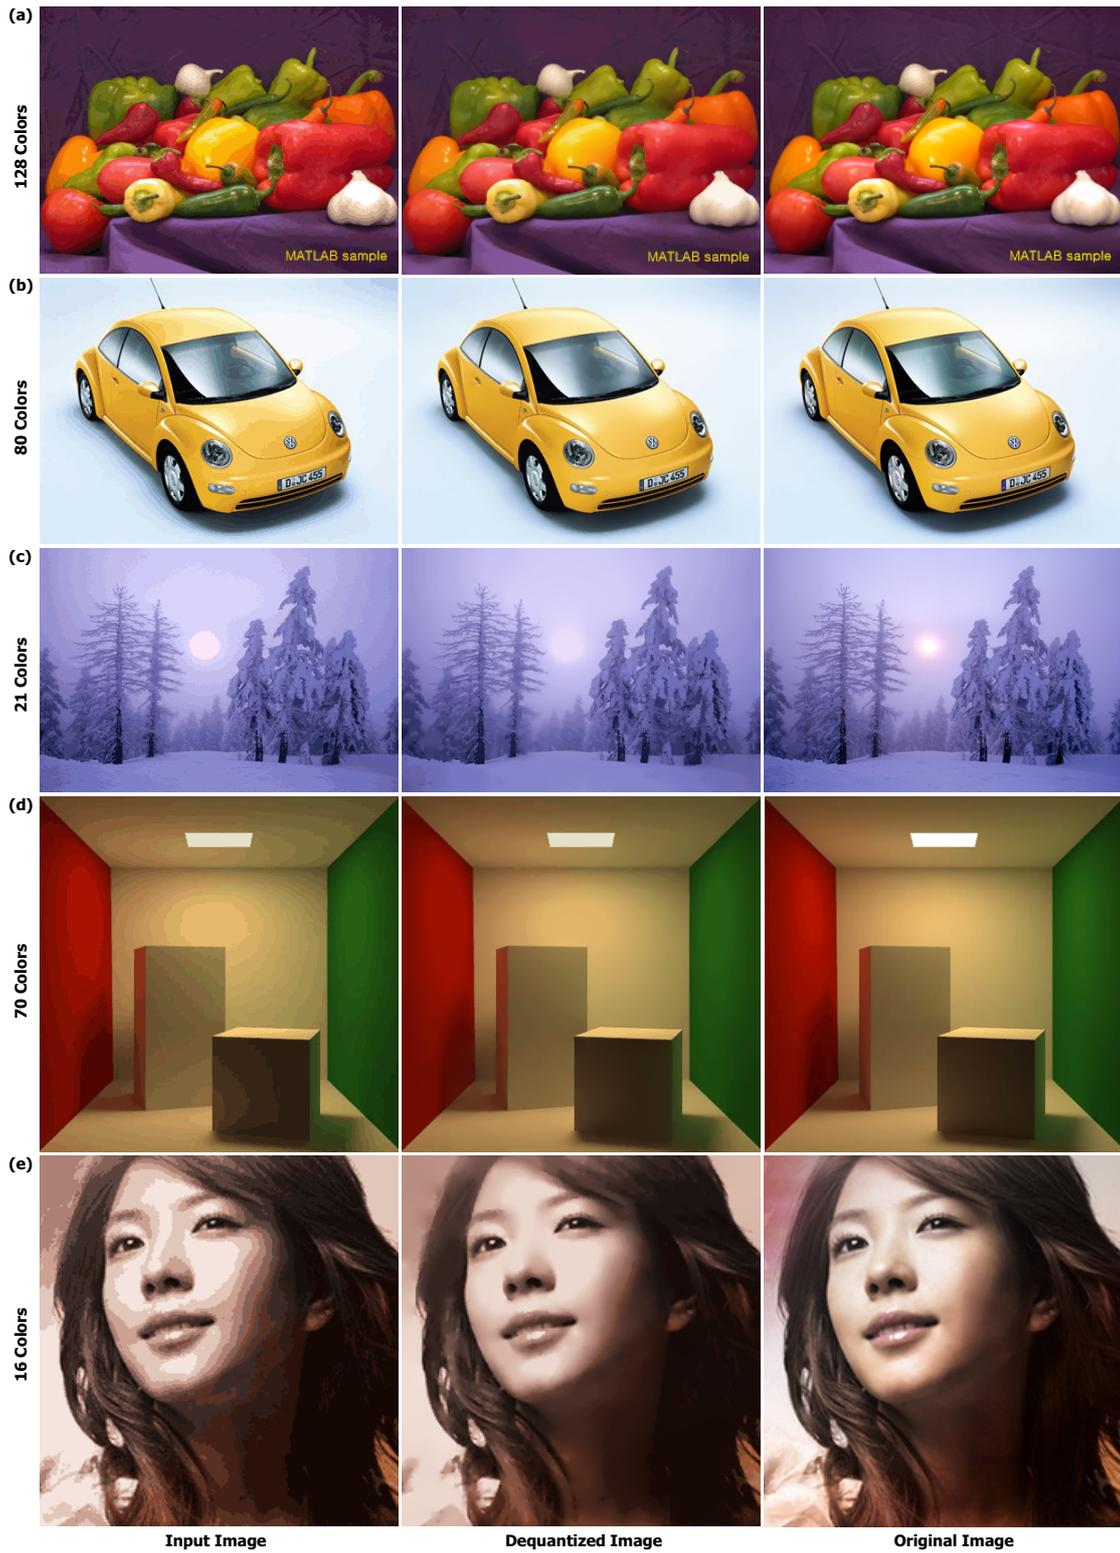


Figure 7: Comparison of color-quantized input, dequantized, and ground truth. The images were dequantized without any user intervention.

Table 1: MSEs of the input image and the dequantized image with respect to the ground-truth in the examples of Figure 7.

Fig.	Input Image	MSE		# of iterations
		Input Image	Dequantized Image	
7(a)	0.0021	0.0015	5	
7(b)	0.0096	0.0073	5	
7(c)	0.0012	0.0012	6	
7(d)	0.0090	0.0088	5	
7(e)	0.0034	0.0031	6	

be discontinuous during the iteration, practically it was not problematic owing to the smoothness term in $J(\phi)$.

In all the experiments, we obtained invertible dequantized images within 2 to 6 iterations. A large weight ω for the data term in Equation (1) reduces the number of pixels out of their Voronoi cells more quickly as illustrated in Figure 5(a). However, a too large ω may cause artifacts since it makes the optimization devote to preserve the representative colors. On the other hand, a too small ω would like to make the optimization devote to enforce smooth variation of color without preserving the representative colors.

We used the mean squared error (MSE) to evaluate the difference between the dequantized image and the ground truth. Figure 5(b) shows the plot of the mean squared error versus the log-scaled weight. Interestingly, the visual quality of the image dequantized with $\omega = 0.5$, which gives the minimum MSE, was poorer than that of the image with a smaller $\omega = 0.01$ (See Figure 6). This is because the human visual system is more sensitive to discontinuity in the smooth parts of an image [KLLH96] and MSE does not reflect it.

Figure 7 shows images dequantized with iterative refining. Table 1 summarizes the mean squared errors with respect to the ground-truth images. In this experiment, we attempted to enhance the visual qualities of the dequantized images so that we applied $\omega = 0.01$. (Although our scheme could produce images with smaller mean squared errors, their visual qualities were not better than the images in Figure 7.)

2.4. 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 can be situations where it is desirable to override some adjacency relationships interactively. For example, in Figure 8, 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 *optional* 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.

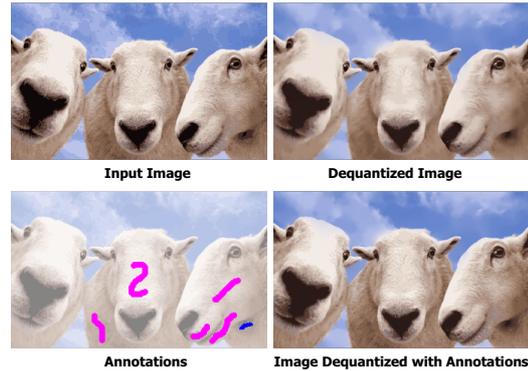


Figure 8: Interactive dequantization with annotations.

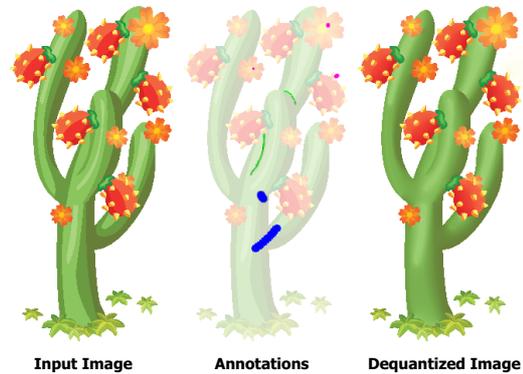


Figure 9: A cartoon-shaded image is dequantized interactively.

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 9 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 10.

2.5. 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 11, the individual smoothness term within each region may produce an unsatisfactory result. The Mach band effect is observed inside the dequantized teapot. To reduce

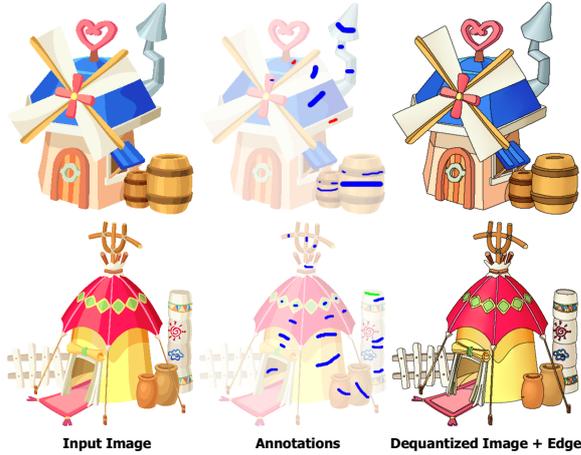


Figure 10: Hard edges are drawn for cartoon shading.

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 $\mathbf{v}_{\mathbf{p},\mathbf{q}}$ 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 $\mathbf{v}_{\mathbf{p},\mathbf{q}}$ should be the same:

$$\varphi(\mathbf{p}) + \varphi(\mathbf{q}) = 2\varphi^*(\mathbf{v}_{\mathbf{p},\mathbf{q}}). \quad (5)$$

In the same vein, we also add continuity constraints along the y -direction. These additional constraints can be added to the objective function (1) as follows:

$$J'(\varphi) = J(\varphi) + w_c \sum_{\substack{\mathbf{p} \in R_i \\ \mathbf{q} \in N_p \cap R_j}} (\varphi(\mathbf{p}) + \varphi(\mathbf{q}) - 2\varphi^*(\mathbf{v}_{\mathbf{p},\mathbf{q}}))^2, \quad (6)$$

where R_j is a region neighboring to R_i of which edge on $\mathbf{v}_{\mathbf{p},\mathbf{q}}$ is soft. By adjusting the weighting factor w_c , we can control the significance of continuity at boundaries, as illustrated in Figure 11.

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 boundaries into soft and hard. Together with this classification, an optimization technique 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 and making pixels preserve their representative colors. The user can optionally override the classification result by scribbling brushes across the regions. We

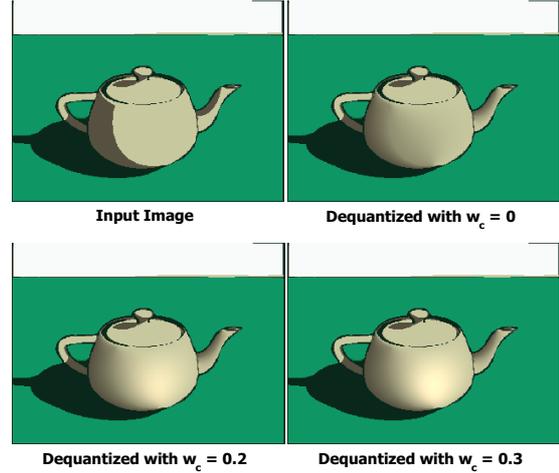


Figure 11: The Mach band illusion ($w = 0$) disappears through the use of additional continuity constraints ($w \neq 0$).

have demonstrated the effectiveness of our approach through diverse examples, such as photographs, cartoons, and artistic illustrations.

Acknowledgments

This work was supported in part by Seoul R&BD Program, Ministry of Information and Communication under the Information Technology Research Center (ITRC) support program, and the Research Grant of Kwangwoon University in 2006.

References

- [BEN04] BEN-EZRA M., NAYAR S. K.: Motion-based motion deblurring. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 26, 6 (2004), 689–698.
- [BK00] BAKER S., KANADE T.: Hallucinating faces. In *Proc. the Fourth International Conference on Automatic Face and Gesture Recognition* (Grenoble, France, 2000).
- [BS98] BORMAN S., STEVENSON R.: *Spatial Resolution Enhancement of Low-resolution Image Sequences - A Comprehensive Review with Directions for Future Research*. Technical report, Laboratory for Image and Signal Analysis (LISA), University of Notre Dame, 1998.
- [BSCB00] BERTALMIO M., SAPIRO G., CASELLES V., BALLESTER C.: Image inpainting. In *Proc. ACM SIGGRAPH* (2000), pp. 417–424.
- [DCOY03] DRORI I., COHEN-OR D., YESHURUN H.: Fragment-based image completion. *ACM Transactions on Graphics* 22, 3 (2003), 303–312.
- [FC02] FUNG Y.-H., CHAN Y.-H.: An iterative algorithm

- for restoring color-quantized images. In *Proc. International Conference on Image Processing* (2002), pp. 313–316.
- [FJP02] FREEMAN W. T., JONES T. R., PASZTOR E. C.: Example-based super-resolution. *IEEE Computer Graphics and Applications* 22, 2 (2002), 56–65.
- [GC91] GALATSANOS N. P., CHIN R. T.: Restoration of color images by multichannel Kalman filtering. *IEEE Transactions on Signal Processing* 39, 10 (1991), 2237–2252.
- [GGS*04] GORELICK L., GALUN M., SHARON E., BASRI R., BRANDT A.: Shape representation and classification using the poisson equation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition* (2004), pp. 61–67.
- [KLLH96] KIM K. M., LEE C. S., LEE E. J., HA Y. H.: Color image quantization and dithering method based on human visual system characteristics. *Journal of Imaging Science and Technology* 40, 6 (1996), 502–509.
- [LLW04] LEVIN A., LISCHINSKI D., WEISS Y.: Colorization using optimization. *ACM Transactions on Graphics* 23, 3 (2004), 689–694.
- [LSA05] LI Y., SHARAN L., ADELSON E. H.: Compressing and companding high dynamic range images with sub-band architectures. *ACM Transactions on Graphics* 24, 3 (2005), 836–844.
- [PGB03] PÉREZ P., GANGNET M., BLAKE A.: Poisson image editing. *ACM Transactions on Graphics* 22, 3 (2003), 313–318.
- [SCI05] SHECHTMAN E., CASPI Y., IRANI M.: Space-time super-resolution. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 27, 4 (2005), 531–545.
- [VNdW*03] VILLE D. V. D., NACHTEGAEL M., DER WEKEN D. V., KERRE E. E., PHILIPS W., LEMAHIEU I.: Noise reduction by fuzzy image filtering. *IEEE Transactions on Fuzzy Systems* 11, 4 (2003), 429–436.
- [WAM02] WELSH T., ASHIKHMIN M., MUELLER K.: Transferring color to greyscale images. *ACM Transactions on Graphics* 21, 3 (2002), 277–280.
- [WOS05] WEISSMAN T., ORDENTLICH E., SEROUSSI G.: Universal discrete denoising: Known channel. *IEEE Transactions on Information Theory* 51, 1 (2005), 5–28.