Removing Quantization Artifacts in Color Images Using Bounded Interval Regularization

Tom Ouyang* and Jack Tumblin**

Northwestern University

Fig. 1. (left) Input image quantized to 16 levels/color input image that shows visible contouring artifacts (right) restored using the method proposed in this paper.

Abstract. Coarsely quantized images will exhibit false contours in smooth low gradient regions. Images intended for standard displays such as CRT monitors can show contours when moved to high dynamic range devices such as HDR displays and film. While various methods exist for noise removal and image restoration, they are not able to remove these contouring artifacts completely and can impose substantial blurring. Our method performs iterative regularization within bounded intervals to remove false contours while preserving natural image features.

1 Introduction

Whenever a scene composed of continuous intensity values is quantized and stored in a digital format, such as when using a digital camera or a scanner, there is inevitable distortion and data loss. Quantization divides the range of input values into a finite number (Q) of non-overlapping ”quantization levels” and all input values within a given interval q are assigned the same digital value. For example, most digital images are stored with Q = 255 possible values for the intensity at each pixel. When an unquantized

* e-mail:ouyang@northwestern.edu

** e-mail:jet@cs.northwestern.edu

image with intensity ranging from $0 \leq u_{ideal} \leq 1.0$ is quantized linearly, all values from the original image are mapped to discrete integer values $0 \leq q \leq 255$ in the quantized image.

The noticeable banding in Figure 1, shown with Q=16 per color channel, is a result of coarse color quantization. These bands occur when smooth regions in the original image are quantized, resulting in abrupt and isolated steps or "contours" in the output image. Pixels on one side of the step are assigned to one quantization interval while pixels on the other side are assigned to an adjacent interval. We refer to this effect as contouring.

Fortunately, most images available today are stored with enough quantization levels Q that visible contouring is rare on ordinary CRT or LCD displays. Even though these images usually do not show quantizing artifacts on standard displays, contouring can become visible when they are used in High Dynamic Range (HDR) displays or in cinema and film applications. As higher contrast displays become more common in professional and consumer markets, this problem will become more prominent. This paper proposes a novel method that repairs contouring due to coarse quantization in digital images.

2 Motivation

Repairing quantization errors falls under the broader category of image restoration and reconstruction methods, which attempt to recover an estimate $u(x,y)$ of the original ideal image $u_{ideal}(x,y)$ from a distorted image $z(x,y)$, which we will refer to as u, u_{ideal}, and z respectively.

Anisotropic Diffusion Filtering [4] is an iterative method for reducing noise levels while retaining important image information. It relies on measurements of the intensity changes to determine the degree of smoothing to be applied for the given region. Any intensity gradient above a given threshold is considered an edge that exists in the ideal image and is preserved or even enhanced. Given enough iterations, this methods generates solutions that may be very smooth, but can introduce unwanted blurring and loss of image features even when an appropriate edge threshold is chosen, as shown in Figure 2. Since sharp edges are preserved by the gradient threshold, this over-smoothing occurs in regions of gradual intensity change.

Regularization methods in image restoration use prior information about the ideal image (e.g. assuming piecewise or global smoothness) to iteratively improve their estimate of the solution. They also restrict their search to solutions to that closely match the observed input z. These techniques generate an estimate for the ideal image u by minimizing the cost functional E.

$$E = \int \lambda R(u) + \frac{1}{2}(u - z)^2 dxdy \tag{1}$$

Fig. 2. The result from applying (left) edge preserving regularization and (right) anisotropic diffusion to Figure 1 (a).

The first term is $R(u)$, known as the regularization cost functional, measures how well the estimate u satisfies the prior knowledge about the ideal image. The second term $(u-z)^2$ measures the deviation from the original image. The constant λ controls the balance between the regularization cost and the deviation cost. A common choice for $R(u)$ is $|\nabla u|^2$ [7], which penalizes roughness and favors images that are globally smooth. Edge Preserving Regularization [3] [5] builds upon this approach using a regularization functional that discourages smoothing across sharp features. We test this method by modifying the Tikhonov regularization functional to treat step gradients larger than one quantization interval as edges that should be preserved. We will describe this choice in more detail in the following section.

Although Edge Preserving Regularization effectively preserves sharp image features, it is not as effective in removing false contours. In existing regularization methods, the deviation term $(u-z)^2$ penalizes any deviation from the observed data. While this penalty works very well for random noise, it tends to preserve contouring artifacts due to coarse quantization. Contouring persists in low-gradient regions of an image where there is relatively little contribution from the regularization function $R(u)$. In these regions the deviation penalty $(u-z)^2$ dominates the cost functional, resulting in solutions that are excessively faithful to the observed input and fail to remove visible contouring, as shown in Figure 2.

3 Proposed Method

As mentioned above, existing regularization methods penalizes any deviation from the input data when searching for a solution. However, quantization is not an assignment of a signal to a single value but rather as an assignment to an interval. Each value z in the quantized image actually represents a range of possible intensity values that compose the quantization interval q_i. The method proposed in this paper treats values in the observed image as quantization intervals instead of as single measurements, as

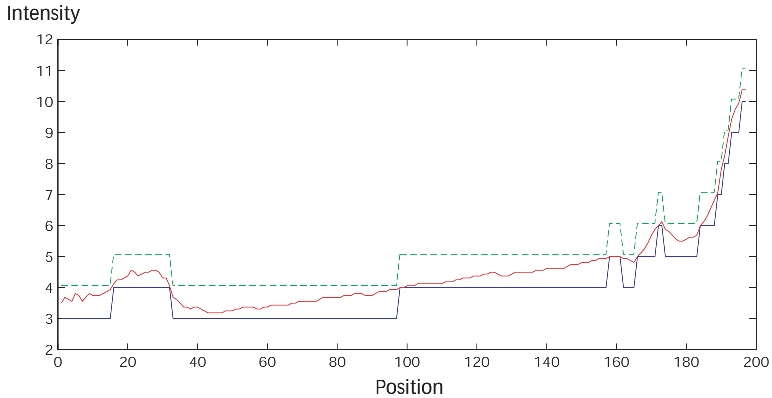

Fig. 3. A scanline representing an ideal unquantized image (red) that is coarsely quantized to the Q=16 levels (blue). The b_{i+1} maximum boundaries for the quantization level are also shown.

show in Figure 3. Consequently, our method only penalizes deviations from the original input when the they exceed the interval boundaries (b_i, b_{i+1}), constraining any solution to stay within the quantization interval q_i. This approach allows us to repair contours from low gradient regions yet prevents over-smoothing of sharp edges.

In order to incorporate the above modifications into our solution, we modify the penalty cost $(u-z)^2$ in Equation 1. We choose the new deviation penalty functional $g(u,b)$ such that it has the following properties.

1. For each pixel in our estimate u of the ideal image, there is no deviation penalty as long as u stays within the original quantization interval boundaries. If the pixel value u belongs to quantization interval q_i, then b_i and b_{i+1} are the lower and upper boundaries for that level.

2. If the estimate u exceeds (b_i, b_{i+1}) boundaries, we impose a strict penalty cost to force u back within the interval:

$$E = \int \lambda R(u) + \frac{1}{2} g(u,b) dxdy \tag{2}$$

where

$$g(u,b) = \begin{cases} 0, & \text{if } b_i <= u <= b_{i+1} \\ (u-b_i)^2, & \text{if } u < b_i \\ (u-b_{i+1})^2, & \text{if } u > b_{i+1} \end{cases} \tag{3}$$

Contouring due to coarse quantization always causes step gradients of one quantization level in the quantized image z. Similarly, any step gradient larger than one quantization level should represent a feature in the unquantized image u_{ideal} that should be preserved. Thus, we can modify our regularization functional $R(u)$ to incur no cost for intensity gradients that exceed one quantization step. We chose a simple implementation in our algorithm by setting $R(u) = 0$ for all pixels u that have a step gradient larger than one quantization interval. Since our method is already constrained from over-smoothing by the interval penalty $g(u,b)$, adding this edge preserving modification provides only limited visible improvements in our tests.

In addition to uniform quantization where each quantization interval is the same size, we also consider non-uniform quantization where intervals have different sizes. Gamma correction is one common source of non-uniform quantization, where more quantization levels are dedicated to the lower (darker) intensities and fewer to the higher (brighter) ones. Our method handles non-uniform quantization by varying values for b_i to match the quantization intervals.

4 Implementation

We use a iterative gradient technique to minimize the cost functional E from Equation (2). We chose the discrete membrane model in Equation (4) from [6] and the deviation penalty functional $g(u,b)$ defined in the previous section. Based on our experiments, the method was able to converge within 10 iterations for a 512x512 color image.

$$E = \sum \lambda((u_{x+1,y} - u_{x,y})^2 + (u_{x,y+1} - u_{x,y})^2) + \frac{1}{2} g(u_{x,y}, b_{x,y}) \tag{4}$$

For each pixel intensity in the current nth iteration $u^n_{x,y}$, we use the discrete update algorithm below to determine the intensity in the next iteration $u^{(n+1)}_{x,y}$:

If $b_i(x,y) <= u_{x,y} <= b_{i+1}(x,y)$,

$$u^{n+1}_{x,y} = u^n_{x,y} - \omega\{4u^n_{x,y} - (u^n_{x-1,y} + u^n_{x+1,y} + u^n_{x,y-1} + u^n_{x,y+1})\}$$

otherwise,

$$u^{(n+1)}_{x,y} = u^n_{x,y} - \omega\{(1+4\lambda)u^n_{x,y} - b_k(x,y) - \lambda(u^n_{x-1,y} + u^n_{x+1,y} + u^n_{x,y-1} + u^n_{x,y+1})\}$$

where $b_i(x,y)$ and $b_{i+1}(x,y)$ are the lower and upper quantization interval boundaries for the pixel $u_{x,y}$, $\omega = 0.5$, $\lambda = 1.0$, and

$$b_k(x,y) = \begin{cases} b_i(x,y), & \text{if } u_{x,y} < b_i(x,y) \\ b_{i+1}(x,y), & \text{if } u_{x,y} > b_{i+1}(x,y) \end{cases} \tag{5}$$

5 Results

We found significant improvement in both contour removal and edge preservation in the restored images using our method. Figure 4 provides a scan-line comparison between our method, edge preserving regularization, and anisotropic filtering. The figure shows that our method is able to generate a smooth estimate of the ideal image that stays completely inside the quantization boundaries. However, edge preserving regularization was not able to completely remove the contours in the image, which still displays partially smoothed contours. There are also areas where the solution deviates from the quantization boundaries and is inconsistent with the ideal image. While result of anisotropic diffusion was a smoother curve with less significant contours, there are significant deviations from the original input. Since the gradient on the right side of the scan-lines is not sharp enough to trigger the edge threshold, both methods over-smoothed that region.

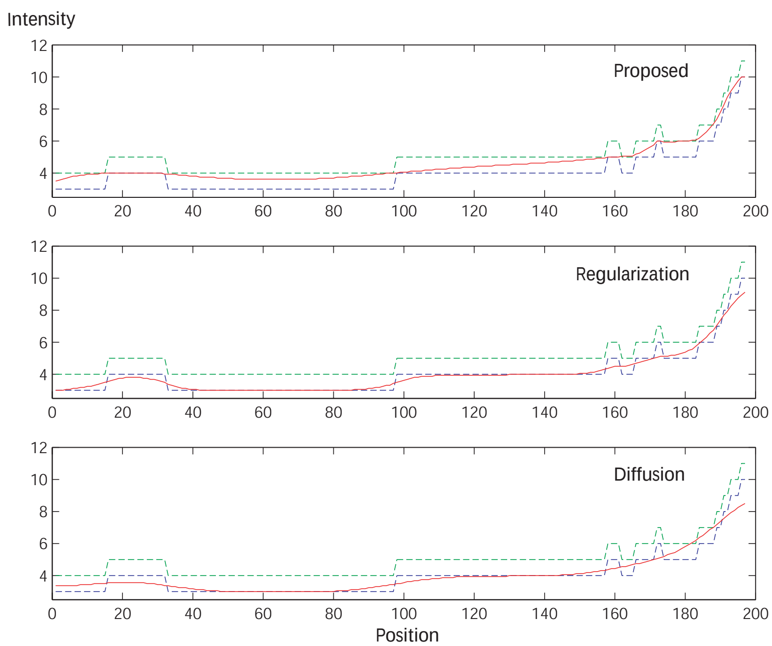

Fig. 4. The scan-line graph from solution generated by (top) the proposed method, (middle) edge-preserving regularization, and (bottom) anisotropic diffusion. The quantization interval boundaries are also included for reference.

In figures 5, 6, and 7 our method corrects coarsely quantized images with very visible contouring artifacts by removing the contours in the sky while maintaining detail in the rest of the image. Figure 5(c) shows the result from existing regularization method, which was not able to fully repair the quantization artifacts and causes more blurring.

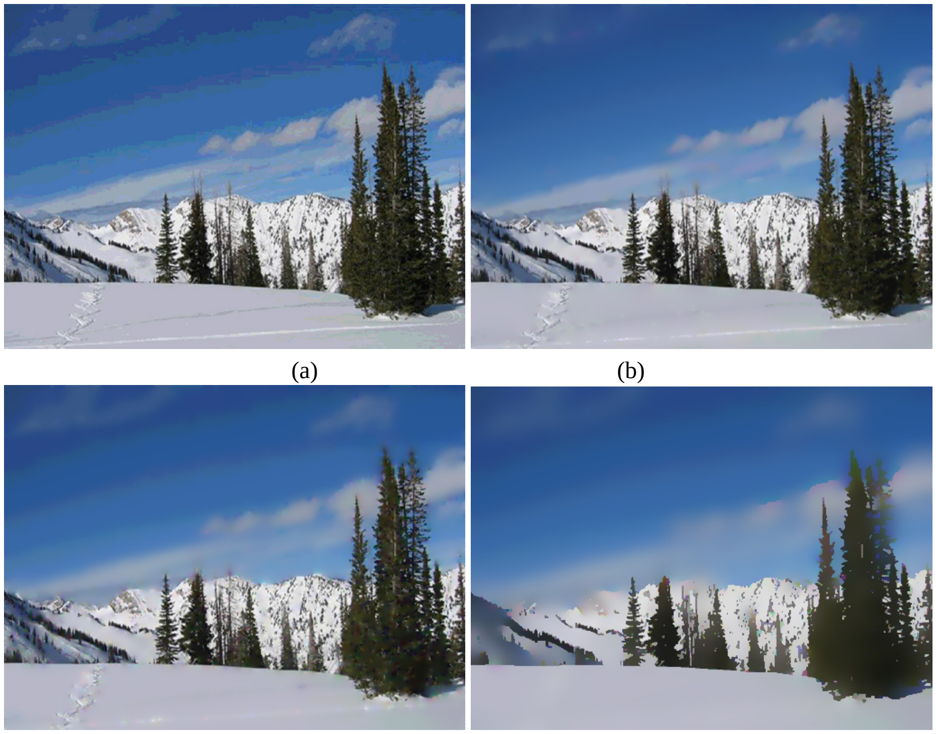

(a) (b) (c) (d)

Fig. 5. Reconstruction of a coarsely quantized image. (a) Distorted Q=16 image restored to 8-bits/color using (b) our method, (c) edge preserving regularization, and using (d) anisotropic diffusion.

6 Related Work

A related approach described in [2] uses a POCS based iterative method for restoring color quantized images. This work addresses the more general problem of restoring an image that has been quantized to a limited color palettes. As a result, the process is much more complex and involves estimating variance of the quantized image and the ideal image for a given color palette using a large training set and requantizing the estimate at every iteration. In addition, they use a low pass filtering constraint that may require many iterations to remove wide contours. In contrast, our method attempts to remove false contours from a coarsely quantized image.

Another method suggested for restoring quantized images [1] converts the image to a level set representation and interpolate from level set contours back to pixels. While this process generates a smooth surface between these contours, it does not ensure that the interpolated solution is consistent with the quantization interval boundaries. Instead of identifying and interpolating isocontours, our method iteratively relaxes the solution within the quantization interval boundaries.

Block reduction is a related method used to repair artifacts from compressed images [8]. Transform coding compression techniques such as JPEG divide the image into small blocks, transform these block to the frequency domain, and store the quantized coefficients. Block reduction techniques attempt to repair artifacts that can appear between neighboring blocks. Similar to our approach, these techniques also impose a constraint based on prior knowledge of the quantization bounds. However, these constraints are based on the quantization of the local transform coefficients in the frequency domain while our method uses the quantization of color intensities to guide image restoration.

7 Discussion

This paper develops a novel regularization approach to repair contouring artifacts on color quantized images. Our approach treats the intensities of the input image as an interval instead of individual measurements. This approach allows our method to remove contouring artifacts that are preserved by other methods.

The proposed method is most effective for images where the only significant error is due to quantization and is not intended as a general noise removal algorithm. For noisy images, it is possible to relax the deviation penalty $g(u)$ from (3) to handle to noise. It may also be possible apply another algorithm such as Total Variation regularization to first remove the noise before using our method to repair any contouring artifacts.

References

1. V. Caselles, J-M. Morel, and C. Sbert. An axiomatic approach to image interpolation. volume 7, pages 376–379.
2. Y.-H. Fung and Y.-H. Chan. Pocs-based algorithm for restoring colour-quantised images. In *IEEE Proceedings - Vision, Image and Signal Processing*, volume 151, pages 119–127, 2004.
3. L. Blanc-Feraud G. Aubert, M. Barlaud and P. Charbonnier. Deterministic edge-preserving regularization in computed imaging. In *IEEE Transactions on Image Process*.
4. P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on PAMI*, 12(7):629–639, 1990.
5. S. Teboul, L. Blanc-Feraud, G. Aubert, and M. Barlaud. Variational approach for edge-preserving regularization using coupled pdes. In *IEEE Transactions on Image Process*, volume 7, pages 387–397, 1998.
6. Demetri Terzopoulos. Regularization of inverse visual problems involving discontinuities. *IEEE Trans. Pattern Anal. Mach. Intell.*, 8(4):413–242, 1986.
7. A. N. Tikhonov. Regularization of incorrectly posed problems. *Sov. Math. Dokl*, 4(1):1624–1627, 1963.
8. A. Zakhor. Iterative procedures for reduction of blocking effects in transform image coding. In *IEEE Transactions on Circuits and Systems for Video Technology*, volume 2, 1992.

Fig. 6. Reconstruction of a coarsely quantized image. (left) Coarsely quantized image with Q=16 (right) restored to Q=255 using the proposed method

Fig. 7. Reconstruction of a coarsely quantized image. (left) Coarsely quantized image with Q=16 (right) restored to Q=255 using the proposed method