Removing Quantization Artifacts In Color Images Using Bounded Interval Regularization

Tom Yu Ouyang
DEPARTMENT OF COMPUTER SCIENCE

Jack Tumblin
FACULTY ADVISOR, DEPARTMENT OF COMPUTER SCIENCE, NORTHWESTERN UNIVERSITY

ABSTRACT

Coarsely quantized images can exhibit false contours in smooth low gradient regions. Images intended for standard displays such as cathode ray tube (CRT) monitors can show contours when moved to high dynamic range (HDR) devices such as HDR displays and film. While various other methods such as regularization and anisotropic diffusion 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.

I. 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 are assigned the same value. For example, most digital images are stored with Q = 255 possible values for the intensity at each pixel. When an unquantized image with intensity ranging from 0 to 1.0 is quantized linearly, all values from the original image are mapped to discrete integer values 0-255 in the quantized image.

Figure 1. A coarsely quantized image with Q = 16.

The noticeable banding in the sky in Figure 1 is a result of coarse color quantization, shown with Q = 16 per color channel. Bands sometimes occur where there are large regions of gradual intensity change in the original unquantized image. When these smooth regions are quantized, the resulting visual effect can be an abrupt and isolated step or "contour" in the displayed image, where pixels on one side of the step are assigned to one quantization level and pixels on the other side are assigned to a neighboring level. We refer to this effect as contouring.

Fortunately, most images available today are stored with enough quantization levels that visible contouring is rare on ordinary CRT or liquid crystal displays (LCD). While these images usually do not reveal quantizing artifacts on standard displays, 8-bit contouring can appear on 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 quantization in digital images.

II. MOTIVATION

Repairing quantization errors falls under the broader category of image restoration and reconstruction methods, which attempt to recover an undistorted ideal image $u(x, y)$ from a distorted observed image $z(x, y)$, where (x, y) represent bit coordinates in the image, which we will refer to simply as u and z. Anisotropic Diffusion Filtering[1] 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 smooth, but introduces unwanted blurring and loss of image features even when an appropriate edge threshold is chosen, as shown in Figure 2a. Since sharp edges are preserved by the gradient threshold, this over-smoothing occurs in regions of gradual intensity change.

Regularization methods approach image restoration by using prior information about the ideal image (i.e. piecewise or global smoothness) to iteratively improve their estimate of the solution. They also restrict their search to solutions which closely match the observed input z. Tikhonov Regularization[2, 3] generates an estimate for the ideal image u by minimizing the cost functional E:

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

The first term $R(u)$, also know as the regularization 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 observed image. The constant λ controls the balance between the regularization cost and the deviation cost. A common choice for $R(u)$ is $|\nabla u|^2$, which penalizes roughness and favors images that are globally smooth. One improvement to this method is Edge Preserving Regularization,[4] which modifies the regularization functional to prevent it from smoothing across sharp edges that should be preserved. Another regularization method, Total Variation[5] uses $|\nabla u|$ as the regularization functional and usually performs better than the quadratic functional at preserving edges.

A feature of regularization methods is that the cost functional E penalizes any deviation from the observed data. While this penalty works very well for random noise, it can pose a problem when dealing with contouring due to quantization.

(a) Anisotropic Diffusion (b) Edge Preserving Regularization

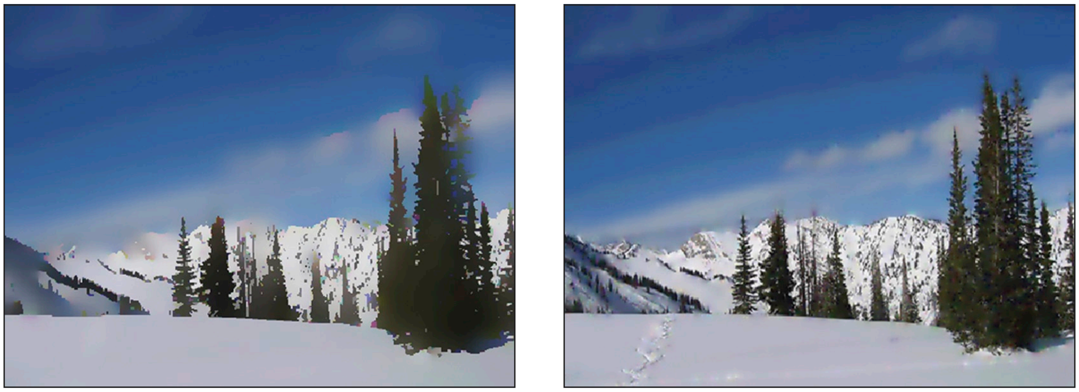

Figure 2. The result from applying **(a)** anisotropic diffusion and **(b)** edge preserving regularization to Figure 1.

Contouring consists of large bands of the same intensity value separated by abrupt intensity gradients. Since the quantization process discarded all changes inside these bands, the roughness cost is very low and the deviation penalty $(u - z)^2$ dominates the search for the solution. This results in images that are very faithful to the observed input but fail to remove much of the visible contouring, as shown in Figure 2b, even when the regularization process is run for a large number of iterations.

III. PROPOSED METHOD

As mentioned above, Tikhonov regularization penalizes any deviation from the initial observed data when searching for a solution. However, each value stored in a quantized image actually represents a range of possible intensity values that compose the quantization interval. The foundation method proposed by these authors builds upon this idea by treating values in the observed image as a quantization interval instead of a single measurement, as show in Figure 3. Our method only penalizes deviations from the observed input when the resulting intensity exceeds the boundaries of the quantization interval. This greatly relaxes the deviation penalty and permits smoothing only as long as the estimated intensities stay within the quantization level boundaries. Because we know that the ideal unquantized intensity must lie within these boundaries, the method places

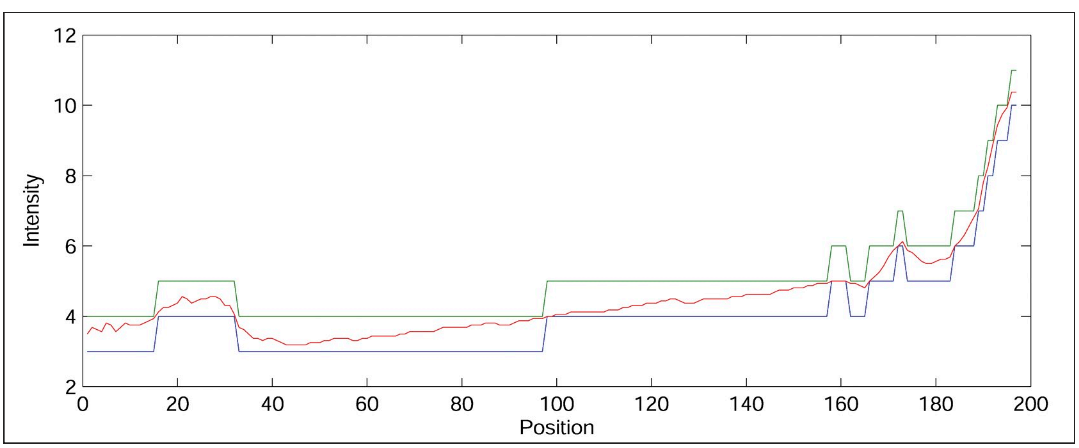

Figure 3. A scan-line representing an ideal unquantized image (red) that is coarsely quantized to the Q=16 levels (blue). The q_{i+1} maximum boundaries for the quantization level are also shown (green line).

a heavy penalty on any deviation that causes the intensity to exceed the quantization interval. This prevents it from over-smoothing both sharp edges and low gradient regions in the image.

In order to incorporate the above modifications into our solution, we use the penalty term $g(u-z)$. The deviation penalty functional $g(x)$ is chosen such that it has the following properties.

1. For each pixel in the estimate of the ideal image, there should be no deviation penalty as long as it stays within the original quantization level boundaries, i.e. the pixel belongs to level i, then q_i and q_{i+1} are the lower and upper boundaries for that level.
2. If the estimate exceeds (q_i, q_{i+1}) boundaries, a strict penalty cost is imposed to force the intensity back within the interval, given by

$$E = \int \lambda R(u) + g(u)\, du \tag{2}$$

where

$$g(u) = \begin{cases} 0 & \text{if} \quad q_i \le u \le q_{i+1} \\ (u-q_{i+1})^2 & \text{if} \quad u > q_{i+1} \\ (u-q_i)^2 & \text{if} \quad u < q_i \end{cases}$$

From Figure 3 we can see that contouring in a coarsely quantized image can occur when two neighboring pixel intensities in a smooth region are assigned to different quantization levels. Because it is not possible for a uniformly smooth low-gradient region in the unquantized image u to become an edge that is larger than one quantization level, we can conclude that deviations >> 1 quantization interval were caused by features in the original image u. Based on this approach, our method ignores changes in intensity that exceed one quantization step in order to preserve the edge. However, since our smoothing is already limited by our deviation penalty, the addition of an edge preserving smoothing term provides only some limited visible improvements.

In addition to uniform quantization where each quantization level is the same size, we also need to consider non-uniform quantization where such levels may not be equal. One common source of non-uniform quantization is gamma correction, 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 q_i to match the quantization scheme.

IV. IMPLEMENTATION

To minimize the cost functional E from Equation 2, our method uses an iterative gradient descent technique. We use the discrete membrane model in Equation 3 proposed by Blake and Zisserman[6] combined with the deviation penalty functional $g(u)$ defined in the previous section. In order to improve the convergence rate our implementation performs Successive Over Relaxation, overcorrecting by a factor ω every iteration. Based on our experiments, the method was able to converge within 20 iterations for a 512×512 color image.

$$E = \lambda\{(u_{x+1,y}-u_{x,y})^2 + (u_{x,y+1}-u_{x,y})^2\} + g(u_{x,y}) \tag{3}$$

For each pixel intensity in the current iteration $u^n_{x,y}$, we use the gradient of the cost E in Equation 3 to determine the intensity in the next iteration $u^{(n+1)}_{x,y}$.

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

Combining Equations 3 and 4, we arrive at the following implementation for the gradient descent algorithm.
For each pixel $u^n_{x,y}$, if $q_i \le u^n_{x,y} \le q_{i+1}$,

$$u^{(n+1)}_{x,y} = u^n_{x,y} - \omega\{4u^n_{x,y} - (u^n_{x-1,y} + u^n_{x+1,y} + u^n_{x,y-1} + u^n_{x,y+1})\} \tag{5}$$

Otherwise, if $u^n_{x,y} < q_i$, $q_k = q_i$, and if $u^n_{x,y} > q_{i+1}$, $q_k = q_{i+1}$

$$u^{(n+1)}_{x,y} = u^n_{x,y} - \omega\{(1+4\lambda)u^n_{x,y} - q_k - \lambda(u^n_{x-1,y} + u^n_{x+1,y} + u^n_{x,y-1} + u^n_{x,y+1})\} \tag{6}$$

V. RESULTS

The simulation results from the proposed method show a significant improvement in both contour removal and edge preservation in the restored images. Figure 4 provides a scan-line comparison between our method (top), edge preserving regularization (middle), and anisotropic filtering (bottom). The figure shows that our method is able to generate a smooth estimate of the ideal image which stays completely inside the quantization boundaries while edge preserving regularization was not able to completely remove the contours in the image and still shows partially smoothed steps. In both edge preserving regularization and anisotropic diffusion there are areas where the solution deviates from the quantization boundaries which is inconsistent with the ideal image.

While the result of anisotropic diffusion provided 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. Figure 5 shows the result of our method on two coarsely quantized images with very visible contouring artifacts. The reconstruction we produced removes the contours in the sky while maintaining detail in the rest of the image.

Figure 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.

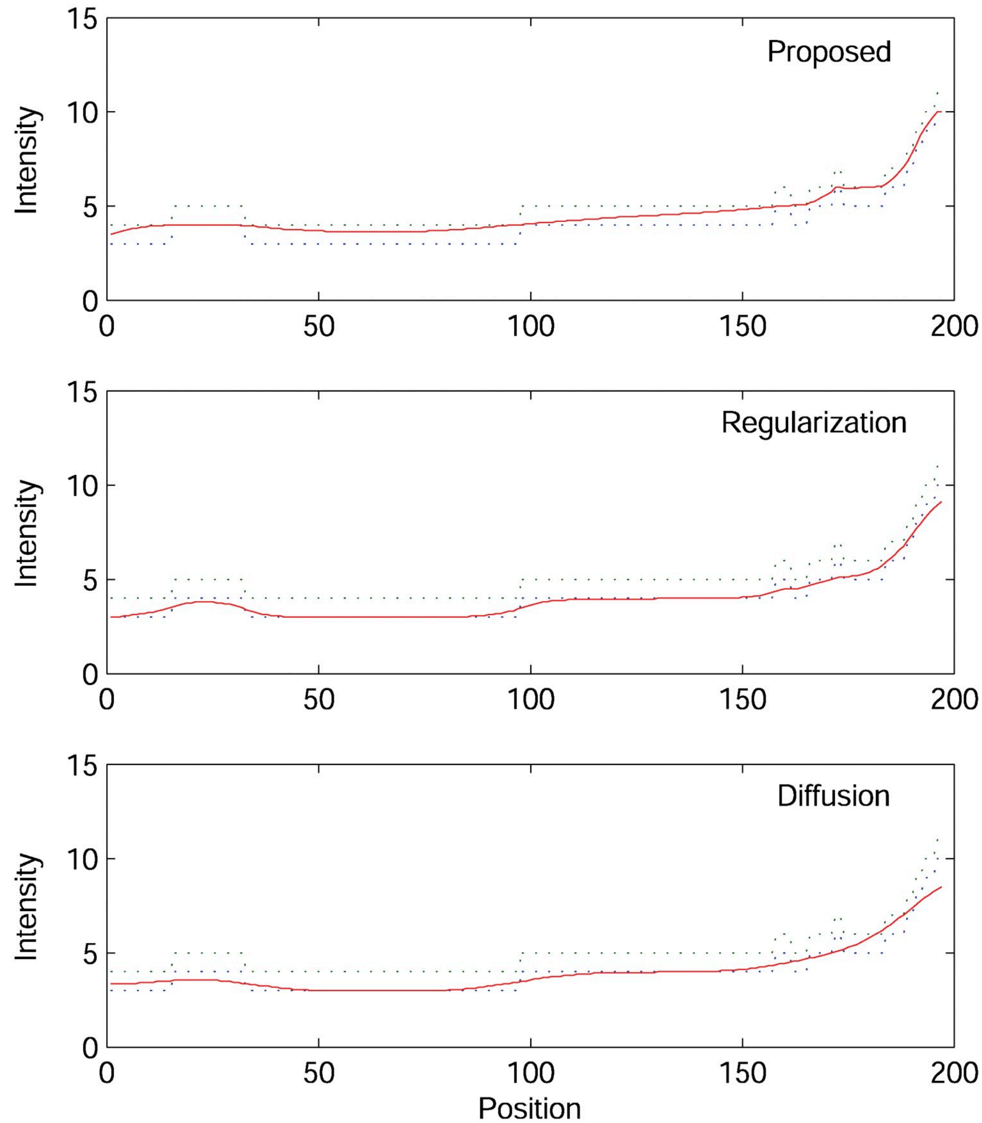

VI. DISCUSSION

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 removal of contouring artifacts that are preserved by other methods.[1,2]

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 Equation 3 to better handle to noise. It may also be possible to apply another algorithm such as Total Variation regularization[5] to first remove the noise before using our method to repair any contouring artifacts.

REFERENCES

1 PERONA, P., AND MALIK, J. 1987. Scale-space and edge detection using anisotropic diffusion. IEEE Trans. PAMI.

2 TIKHONOV, A. N. 1963. Regularization of incorrectly posed problems. Sov. Math. Dokl 4, 1, 1624–1627.

3 RUDIN, L. I., OSHER, S., AND FATEMI, E. 1992. Nonlinear total variation based noise removal algorithms. In Proceedings of the eleventh annual international conference of the Center for Nonlinear Studies on Experimental mathematics : computational issues in nonlinear science, Elsevier North-Holland, Inc., 259–268.

4 G. AUBERT, M. BARLAUD, L. B.-F., AND CHARBONNIER, P. 1997. Deterministic edge-preserving regularization in computed imaging. In IEEE Trans. Imag. Process.

5 VOGEL, C. R., AND OMAN, M. E. 1996. Iterative methods for total variation denoising. SIAM Journal on Scientific Computing 17, 1, 227–238.

6 BLAKE, A., AND ZISSERMAN, A. 1987. Visual reconstruction.MIT Press, Cambridge, M.A..

Figure 5. Reconstruction of coarsely quantized images.
Coarsely quantized image with (**left**) Q=16 restored to (**right**) Q=255 using the proposed method.