

Improving Image Restoration with Soft-Rounding

Xing Mei^{1,3} Honggang Qi² Bao-Gang Hu³ Siwei Lyu¹

¹Computer Science Department, University at Albany, SUNY

²Computer Science Department, University of Chinese Academy of Sciences

³NLPR, Institute of Automation, Chinese Academy of Sciences, Beijing, China

xmei@albany.edu, hgqi@jdl.ac.cn, hubg@nlpr.ia.ac.cn, slyu@albany.edu

Abstract

Several important classes of images such as text, barcode and pattern images have the property that pixels can only take a distinct subset of values. This knowledge can benefit the restoration of such images, but it has not been widely considered in current restoration methods. In this work, we describe an effective and efficient approach to incorporate the knowledge of distinct pixel values of the pristine images into the general regularized least squares restoration framework. We introduce a new regularizer that attains zero at the designated pixel values and becomes a quadratic penalty function in the intervals between them. When incorporated into the regularized least squares restoration framework, this regularizer leads to a simple and efficient step that resembles and extends the rounding operation, which we term as soft-rounding. We apply the soft-rounding enhanced solution to the restoration of binary text/barcode images and pattern images with multiple distinct pixel values. Experimental results show that soft-rounding enhanced restoration methods achieve significant improvement in both visual quality and quantitative measures (PSNR and SSIM). Furthermore, we show that this regularizer can also benefit the restoration of general natural images.

1. Introduction

The task of image restoration is to recover an image from its noisy and/or blurry observation [3]. Image restoration is an example of ill-posed inverse problems, and unique solutions can only be obtained by introducing proper regularizations and constraints on the clean images. Indeed, it is now generally accepted that the more prior knowledge we have about the properties of the uncorrupted images, the better we can solve the image restoration problem [1, 9].

To date, the most widely used statistical regularities for images are obtained from their band-pass filter responses (such as gradients or wavelets) or mean-removed pixel patches. Such responses have been observed to have heavy-

tailed marginal histograms, on the basis of which regularizers or priors that prefer sparse responses have been used to attain the state-of-the-art restoration performances. Examples include regularizers in the form of the ℓ_0 norm [25, 26] and total variation [20], or prior models based on generalized Laplacian [11, 13] and Gaussian mixtures [23, 28]. In these image restoration methods, pixel values are usually treated as continuous variables.

However, images are captured and stored in a digital format, which poses a general constraint on pixel values. The pixels of an image in a b -bit format can only take integer values from the set $\{0, \dots, 2^b - 1\}$. For some important subclasses of images such as text or pattern images, their pixels can take only a distinct set of the full range integer values (see Fig.1(a) for an example). Furthermore, it is usually possible to obtain these values before we restore such an image. For text and barcode images, the pixel values are usually separated into two distinct classes of foreground pixels and background pixels. For pattern images, image pixels usually take multiple distinct pixel values, which can be estimated from other images of the same type.

Knowing that the pristine image can only take some distinct pixel values provides a useful constraint for image restoration, which can help to suppress visual artifacts and improve the restoration performance. A straightforward approach would be to round the restored pixel values to the nearest distinct pixel values as a post-processing step. However, such a simple solution often does more harm than good, as the rounding step can undo the structures recovered from the restoration step (see Fig.1(c)). On the other hand, except in a few cases [10, 16, 18, 27], such prior information about pixel values seems trivial and is largely overlooked by most existing image restoration algorithms.

In this work, we describe an effective and efficient approach to incorporate the knowledge of distinct pixel values of the pristine images into the popular regularized least squares restoration framework. We introduce a new regularizer of the uncorrupted images in the pixel domain. This regularizer attains zero at the designated pixel values and is




Figure 1: An example demonstrating the effect of our method. (a) An original image with four distinct pixel values 0, 100, 150, 255. (b) Its corruption with a 101×101 blur kernel and 1% white Gaussian noise. (c) Restoration results using the state-of-the-art method based on ℓ_0 regularization of image gradients [18]. The artifacts in the restored image are due to pixel values that are not consistent with the four distinct pixel values. (d) Rounding (nearest neighbor) of the restoration results in (c) to the four distinct pixel values. Though some artifacts in (c) are removed, some structures restored in (c) are also destroyed. (e) Restoration results using our method that combines gradient domain ℓ_0 regularization based restoration with soft-rounding operation. Note that our method achieves significant improvement in comparison with (c) and (d) both visually and quantitatively (PSNR and SSIM).

a quadratic function in the intervals between them. As such, it reflects the preference of the pixel values of the restored image to take the designated distinct pixel values. When incorporated into the regularized least squares restoration framework, this regularizer leads to a simple and efficient solution that resembles the rounding operation, which we term as *soft-rounding*. We apply the soft-rounding enhanced solution to various state-of-the-art image restoration methods, and show that it achieves significant improvement in both visual quality and quantitative measures (PSNR and SSIM) for the restoration of binary images and multi-value pattern images. We also demonstrate its effectiveness for general natural image restoration.

2. Background and Related Work

We formulate the image restoration problem as recovering a pristine image of m pixels, which is vectorized and denoted with an m -dimensional vector \mathbf{x} , from its noisy and/or blurry observation \mathbf{y} . We assume that \mathbf{y} is generated from the convolution of \mathbf{x} with a spatially invariant blurring kernel \mathbf{k} and further contaminated with additive Gaussian noise \mathbf{n} , as:

$$\mathbf{y} = \mathbf{k} \otimes \mathbf{x} + \mathbf{n}, \quad (1)$$

where \otimes denotes the convolution operation [3]¹. The restoration problem can be solved within a *regularized least squares* (RLS) estimation framework, as:

$$\min_{\mathbf{x}} \frac{1}{2} \|\mathbf{y} - K\mathbf{x}\|_2^2 + \lambda_N \Gamma_N(D\mathbf{x}) \quad (2)$$

where K corresponds to the block Toeplitz matrix representation of the convolution kernel \mathbf{k} , D represents a lin-

¹In the current work, we consider the setting that the blurring kernel and noise variance are known, known as *non-blind* restoration. Our method can also be extended to blind restoration [14].

ear transform through which image properties can be better modeled, Γ_N is a regularizer in the transformed domain, and λ_N is an adjustable parameter balancing the contributions of the data fidelity term and the regularization term in the objective function. For image denoising, K reduces to an identity matrix.

The choice of the linear transform D and the regularizer Γ_N is essential for effective restoration. A common approach is to first transform image to band-pass domains (e.g., gradients or wavelets) or use mean-removed patches, then use regularizers in the form of sparsity-encouraging norms, e.g., ℓ_0 [25, 26] and total variation [20], or models capturing statistical properties of images in such domains, e.g., generalized Laplacian [11, 13], Gaussian scale mixtures [23], or general Gaussian mixtures [28].

To date, properties of the pristine images in the pixel domain have been largely overlooked in image restoration but in a few cases [4, 10, 16, 18, 27]. The constraint on pixels to take distinct values has been considered in a few recent works in binary image restoration, where the pixels can take only two values. Zhang introduced a regularization term to penalize pixel values that drift away from the two values [27]. The specific regularizer used there is a fourth-order polynomial and requires iterative numerical approximations on each pixel. Pan *et al.* used a ℓ_0 regularization term if only the zero peak in pixel values is considered [18]. They showed that this ℓ_0 intensity prior helps to identify salient edges in text images and therefore improves blur kernel estimation, but its non-blind restoration performance was not reported in the paper. A potential limitation of the ℓ_0 intensity prior is that it might not be suitable for binary images whose pixel values center around two peaks far away from zero. Another common limitation of these binary image restoration methods is that they can not be easily

extended to handle more than two distinct pixel values.

Some recent methods further incorporate histograms of the pixel values in image restoration. Chen *et al.* studied the intensity histograms of clear document images and used them as reference information for blind deconvolution, leading to improved restored results with reduced visual artifacts [4]. Mei *et al.* recently proposed a framework to enforce marginal histogram constraints in image restoration. They showed that for pattern images, marginal intensity histograms can significantly improve the denoising performance under high noisy levels [16]. Although a pixel histogram provides more information than specifying distinct pixel values in an image, estimating such histograms from degraded images is a non-trivial task [5, 16, 29], which poses a limitation for these methods in practice. In contrast, as we pointed out in Section 1, obtaining distinct pixel values before restoring an image is usually possible for the type of images we focus on in this work.

3. Method

As discussed previously, some subclasses of images have distinct pixel values. When these values are known *a priori*, they provide valuable information to improve the restoration performance. In this section, we describe in detail an approach to incorporate such information into the general RLS restoration framework, Eq.(2). It turns out that the specific solution we obtained has a particularly intuitive interpretation as a “soft-rounding” operation, which adaptively adjusts the restored pixel values to the target values.

3.1. Distinct Pixel Value Regularizer

We first introduce a regularizer to encourage pixel values of the restored image to take the designated distinct pixel values. Formally, assuming we require the restored image should have n distinct pixel values $t_1 < t_2 < \dots < t_n$, represented in a vector $\mathbf{t} = (t_1, t_2, \dots, t_n)$, we define the regularizer as

$$\Gamma_I(\mathbf{x}) = \sum_{i=1}^m \gamma_{\mathbf{t}}(x_i) \quad (3)$$

where m is the number of the image pixels, and $\gamma_{\mathbf{t}}$ measures the intensity deviation on each pixel:

$$\gamma_{\mathbf{t}}(x) = \begin{cases} \frac{1}{2}(t_1 - x) & \text{if } x < t_1 \\ \frac{1}{2}(x - t_j)(t_{j+1} - x) & \text{if } x \in [t_j, t_{j+1}] \\ \frac{1}{2}(x - t_n) & \text{if } x > t_n \end{cases} \quad (4)$$

Fig.2 shows an example of $\gamma_{\mathbf{t}}(x)$ in the range of $[0, 1]$. This function has the following distinct characteristics:

1. $\gamma_{\mathbf{t}}(x) \geq 0$ for $x \in \mathcal{R}$;
2. $\gamma_{\mathbf{t}}(x)$ attains zero at t_j ($j = 1, \dots, n$) and reaches local maximum when x locates at the midpoint of each interval;




Figure 2: An example of $\gamma_{\mathbf{t}}(x)$ in the range of $[0, 1]$, with five distinct values (\diamond) corresponding to $[0.1, 0.5, 0.7, 0.8, 0.9]$.

3. for $x \in [t_j, t_{j+1}]$, $\gamma_{\mathbf{t}}(x)$ is a concave quadratic function, and $\gamma_{\mathbf{t}}(x)$ reduces to a linear function outside the range of $[t_1, t_n]$.

Note that $\gamma_{\mathbf{t}}(x)$ is neither a convex nor a concave function over the whole real line, but at individual intervals it is concave (it becomes either a quadratic or linear function).

3.2. Using Γ_I in Image Restoration

We incorporate regularizer $\Gamma_I(\mathbf{x})$ into the RLS restoration framework and solve the following problem:

$$\min_{\mathbf{x}} \frac{1}{2} \|\mathbf{y} - K\mathbf{x}\|_2^2 + \lambda_N \Gamma_N(D\mathbf{x}) + \lambda_I \Gamma_I(\mathbf{x}) \quad (5)$$

where we introduce two parameters λ_N, λ_I to balance the contributions of the three terms in the overall objective function. Note that Eq.(5) is a non-convex problem, and we optimize it for a local minimum.

Following the general variable splitting scheme [2], we introduce an auxiliary variable \mathbf{z}_I and convert Eq.(5) to the following equivalent problem:

$$\begin{aligned} \min_{\mathbf{x}} \quad & \frac{1}{2} \|\mathbf{y} - K\mathbf{x}\|_2^2 + \lambda_N \Gamma_N(D\mathbf{x}) + \lambda_I \Gamma_I(\mathbf{z}_I) \\ \text{s.t.} \quad & \mathbf{z}_I = \mathbf{x} \end{aligned} \quad (6)$$

We then solve this equality constrained problem using the *augmented Lagrangian method* (ALM) [17]. Introducing a Lagrange multiplier \mathbf{w}_I to the equality constraint and a penalty parameter μ_I , we form the augmented Lagrangian function of problem (6), as

$$\mathcal{L}(\mathbf{x}, \mathbf{z}_I, \mathbf{w}_I) = \frac{1}{2} \|\mathbf{y} - K\mathbf{x}\|_2^2 + \lambda_N \Gamma_N(D\mathbf{x}) + \lambda_I \Gamma_I(\mathbf{z}_I) + \frac{\mu_I}{2} \|\mathbf{x} - \mathbf{z}_I\|_2^2 - \mathbf{w}_I^T (\mathbf{x} - \mathbf{z}_I) \quad (7)$$

Starting with initial values for $\mathbf{z}_I, \mathbf{x}, \mathbf{w}_I$ and parameter μ_I , we iteratively update $\mathbf{z}_I, \mathbf{x}, \mathbf{w}_I$ with regards to the augmented Lagrangian starting from $k = 1$ as follows:

1. Update \mathbf{z}_I : $\mathbf{z}_I^{k+1} = \operatorname{argmin}_{\mathbf{z}_I} \mathcal{L}(\mathbf{x}^k, \mathbf{z}_I, \mathbf{w}_I^k)$
2. Update \mathbf{x} : $\mathbf{x}^{k+1} = \operatorname{argmin}_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \mathbf{z}_I^{k+1}, \mathbf{w}_I^k)$
3. Update \mathbf{w}_I : $\mathbf{w}_I^{k+1} = \mathbf{w}_I^k - \mu_I (\mathbf{x}^{k+1} - \mathbf{z}_I^{k+1})$.

It is guaranteed in theory that the solution to the original problem (6) can be obtained by solving (7) iteratively in the ALM framework when penalty parameter μ_I increases

to a sufficiently large value [17]. Compared with the half-quadratic splitting method [8, 11], a significant advantage of ALM is that its convergence can be assured without increasing μ_I indefinitely, which is especially suitable for solving problems with multiple equality constraints.

The update step for the multiplier (step 3) is straightforward as required by the ALM framework [17]. The subproblem updating \mathbf{x} (step 2) can be reformulated as a RLS problem:

$$\min_{\mathbf{x}} \frac{1}{2} \|\mathbf{y} - K\mathbf{x}\|_2^2 + \frac{\mu_I}{2} \left\| \mathbf{z}_I^{k+1} + \frac{\mathbf{w}_I^k}{\mu_I} - \mathbf{x} \right\|_2^2 + \lambda_N \Gamma_N(D\mathbf{x}) \quad (8)$$

Compare to Eq.(2), Eq.(8) has an additional quadratic term about \mathbf{x} , *i.e.*, the second term. If Γ_N is also a quadratic term on \mathbf{x} , this subproblem can be solved through conjugate gradient or FFT-based method. For most existing restoration methods with non-smooth Γ_N regularizers, this subproblem can be further decomposed into smaller subproblems through half-quadratic splitting or ALM [8, 11], for which well-developed proximal operators for existing Γ_N regularizers can be reused in the computation process. Note that for practical implementations, the decomposition of the \mathbf{x} subproblem can be done within the same ALM framework as \mathbf{z}_I , and its solution requires only a one-iteration approximate updating process. Compared to standard restoration methods with only Γ_N terms, our method requires extra slight computation costs on the \mathbf{z}_I and \mathbf{w}_I updating steps. For better understanding of the \mathbf{x} subproblem, please refer to the supplemental material for two concrete examples involving different Γ_N terms. The solution of the \mathbf{z}_I (step 1) is detailed in the following section.

3.3. Soft-Rounding Operator

After rearranging terms, the subproblem of updating \mathbf{z}_I (step 1) reduces to the following minimization problem:

$$\min_{\mathbf{z}_I} \frac{\mu_I}{2} \|\mathbf{z}_I - \mathbf{c}\|_2^2 + \lambda_I \Gamma_I(\mathbf{z}_I), \quad (9)$$

where we define $\mathbf{c} = \mathbf{x}^k - \frac{1}{\mu_I} \mathbf{w}_I^k$. Note that both terms in Eq.(9) can be split into the sum on each element of \mathbf{z}_I (see the definition of $\Gamma_I(\cdot)$ in Eq.(3)). Therefore this problem can be optimized separately on each element as

$$(\mathbf{z}_I^{k+1})_i = \phi_{\mathbf{t}, \frac{\lambda_I}{\mu_I}}(c_i), \quad (10)$$

where

$$\phi_{\mathbf{t}, \lambda}(c) = \operatorname{argmin}_x \left(\frac{1}{2\lambda} (x - c)^2 + \gamma_{\mathbf{t}}(x) \right) \quad (11)$$

can be seen as the *proximal operator* of function $\gamma_{\mathbf{t}}(\cdot)$ [19].

The following result shows that $\phi_{\mathbf{t}, \lambda}(c)$ can be computed with a piece-wise function.




Figure 3: Examples of function $\phi_{\mathbf{t}, \lambda}(c)$ for different λ values and the same set of $\mathbf{t} = [0.1, 0.5, 0.7, 0.8, 0.85]$ (\diamond). Left: soft-rounding ($\lambda = 0.6$). Right: rounding ($\lambda = 1.1$).

Theorem 1. For constant $\lambda > 0$ and $t_1 < t_2 \cdots < t_n$, the optimal solution to Eq.(11) is given by:

- (i) for $c < t_1$, $\phi_{\mathbf{t}, \lambda}(c) = \min(t_1, c + \frac{\lambda}{2})$;
- (ii) for $c > t_n$, $\phi_{\mathbf{t}, \lambda}(c) = \max(t_n, c - \frac{\lambda}{2})$;
- (iii) hard-rounding: for $c \in [t_1, t_n]$ and $\lambda \geq 1$,

$$\phi_{\mathbf{t}, \lambda}(c) = \begin{cases} t_j & c \in [t_j, t_j + d_j) \\ t_{j+1} & c \in (t_{j+1} - d_j, t_{j+1}] \end{cases} \quad (12)$$

where $d_j = \frac{1}{2}(t_{j+1} - t_j)$ and $j = 1, \dots, n-1$;

- (iv) soft-rounding: for $c \in [t_1, t_n]$ and $\lambda \in (0, 1)$,

$$\phi_{\mathbf{t}, \lambda}(c) = \begin{cases} t_j & c \in [t_j, t_j + d_j] \\ \frac{c}{1-\lambda} - \frac{\lambda(t_j + t_{j+1})}{2(1-\lambda)} & c \in [t_j + d_j, t_{j+1} - d_j] \\ t_{j+1} & c \in [t_{j+1} - d_j, t_{j+1}] \end{cases} \quad (13)$$

where $d_j = \frac{\lambda}{2}(t_{j+1} - t_j)$ and $j = 1, \dots, n-1$.

The proof of Theorem 1 is given in the supplementary material. In Fig.3 we show examples of $\phi_{\mathbf{t}, \lambda}$ with different λ values and the same set of \mathbf{t} . As can be seen from Theorem 1 and Fig.3, when $\lambda \geq 1$, $\phi_{\mathbf{t}, \lambda}$ maps $c \in [t_j, t_{j+1}]$ to the nearest of t_j or t_{j+1} , *i.e.*, it rounds c to either t_j or t_{j+1} . For $0 < \lambda < 1$, $\phi_{\mathbf{t}, \lambda}$ keeps a zone $[t_j + d_j, t_{j+1} - d_j] \subset [t_j, t_{j+1}]$, inside which c is not rounded but undergoes a linear transform, the slope of which is determined by λ . As such, we term Eq.(13) as *soft-rounding*, drawing an analogy from *soft-thresholding*, which is the proximal operator of ℓ_1 norm [7]. The soft-rounding operator “softens” the rounding operator, and has an interesting effect on its output: if c is close to the endpoints of an interval, it will be rounded to the nearest endpoint. However, if c lies in the middle range of an interval, it is “nudged” towards the nearest t value through a linear interpolation.

Note that through the ALM framework, regularizer Γ_I and the resulting soft-rounding operator can readily work with most existing RLS image restoration methods. The two regularizers, Γ_N and Γ_I , collaborate in a closed loop to recover the uncorrupted image. The former restores structures (edges and contours), and the latter encourages the

restored image to have the desirable distinct pixel values. Also, typical restoration methods tend to smooth the image, and the soft-rounding operator, by enhancing the contrasts of the restored image, serves as a countermeasure to such over-smoothing.

Our method is advantageous to the simpler approach that first restores the image and then applies rounding to project the pixel values back to the target values. Since rounding is performed independent of and after the restoration step, it may destroy the structures recovered by restoration and undo its effect. This difference is demonstrated in the subsequent experiments in Section 4.

3.4. Obtaining Distinct Pixel Values

We first discuss several general ways to obtain distinct pixel values to be used in Γ_I , depending on the application scenario and domain knowledge. (1) Distinct pixel values can be extracted from external sources. For instance, text images captured under similar illumination conditions might vary in spatial layout and visual contents, but they show very close distinct values for text and background pixels respectively. (2) Distinct pixel values can be specified through user interaction. We first recover a clean image from the degraded observation with existing restoration methods, and then manually select dominant pixel values from various regions as distinct pixel values. (3) For specific problems, distinct pixel values can be estimated from the degraded observation. For images contaminated with small Gaussian noise, their marginal intensity histograms might be roughly estimated from the noisy observations following a 1D deconvolution approach [29], from which we can extract distinct pixel values.

Apart from general discussions, we also provide a simple but effective method to estimate distinct pixel values from a degraded image. A natural idea is to first recover a clean image with standard image restoration methods (e.g., L0) and then use clustering methods such as K-Means to extract the cluster centers as distinct pixel values. However, the image degradation process blurs across different distinct pixel values and shrinks the distances between them. This effect might still be evident in the restored image, sometimes leading to inaccurate centers as shown Table 1. Inspired by the local two-color model [10] and its recent extension [12], we propose a simple two-step estimation method. In the first step, we extract a small patch around each pixel and estimate a local one-center or two-center model. If the intensity variance in the patch is below a threshold, the patch falls into the one-center model, and its mean pixel value is collected for the second global clustering step; otherwise the patch is a two-center patch, and its two cluster centers are collected with a *normalized* two-class K-Means clustering method [12]. Compared to K-Means, normalized K-Means tries to stretch the distance between the two cluster

Method	Examples		
	Fig. 1	Fig. 4	Fig. 6 (2nd Row)
True Values	0, 100, 150, 255	26, 217	24, 53, 158, 224, 255
K-Means	18, 98, 156, 245	62, 215	14, 55, 164, 221, 253
Our Method	9, 101, 155, 255	32, 217	20, 52, 162, 224, 255

Table 1: The distinct pixel values estimated by K-Means and our method for three L0 deconvolution visual examples.

centers and fight against the shrinkage effect due to image degradation. Then in the second step, we perform a global K-Median clustering on all the collected local center values to determine the final distinct pixel values. We provide the estimation results for three L0 deconvolution text and pattern image examples in Table 1: the pixel values estimated by our method are very close to those of the clean images and are more accurate than the results produced by a single global clustering method such as K-Means.

4. Experiment

In this section, we combine the soft-rounding operator with several state-of-the-art RLS image restoration methods and evaluate its performance on a variety of images. We focus on the restoration of three typical categories of images with distinct pixel values, namely, text images with two pixel values, pattern images with multiple pixel values, and natural images with pixels of full range of 8-bit depth. Furthermore, we also test our algorithm on other types of images that are of practical importance, such as barcodes, posters and license plates. In all experiments, we assume the distinct pixel values are known.

4.1. Text Image Deconvolution

Text images form an important but special class of images. Many text images originate from text documents and have two distinct pixel values for text and background pixels respectively. State-of-the-art non-blind restoration performance for text images was achieved with an RLS method with an ℓ_0 regularizer in the gradient domain (subsequently denoted as L0 in Table 2) [18, 25, 26]. Assuming the pixel values for foreground and background pixels have been extracted from similar text images, we augment the ℓ_0 -based RLS restoration method with soft-rounding and denoted it as L0+S. We also compare the performance of L0 combined with three other pixel domain regularizers: a simple post-processing method which directly rounds the restored image to a binary image (denoted as L0+R), the fourth-order polynomial regularizer to incorporate binary intensity constraints [27] (denoted as L0+P) and the ℓ_0 intensity regularizer from [18] (denoted as L0+L). For fair comparison of each method, we adjust parameters and report the best overall performance.

We collect 20 text images of different languages, equations and graphs. We use three kernel/noise settings to simulate the degradation process: (i) a 33×33 kernel and




Figure 4: Text Image Deconvolution Example. The original image is degraded with a 51×51 kernel and 1% Gaussian noise. The values for text and background pixels are $t_1 = 26, t_2 = 217$ respectively. Details are best viewed on screen.

additive Gaussian noise with 3% standard deviation; (ii) a 45×45 kernel and 2% Gaussian noise; and (iii) a 51×51 kernel and 1% Gaussian noise. The kernels are motion-blur kernels previously used in [13, 18].

Table 2 reports the average peak-signal-to-noise (PSNR) and structural similarity index (SSIM) [24] results of the five methods in three kernel-noise settings. Compared to the original L0 method, the soft-rounding enhanced restoration method significantly improves the average PSNR by at least 0.5dB and the average SSIM by at least 0.04 for all the settings, which is more effective than the other two methods with intensity regularizers (L0+L, L0+P). The simple rounding method (L0+R) actually lowers the performance for most settings. These quantitative results demonstrate the effectiveness of our method.

Kernel/Noise	Avg. PSNR (in dB)				
	L0	L0+R	L0+L	L0+P	L0+S
$33 \times 33+3\%$	19.46	16.75	19.71	19.91	20.03
$45 \times 45+2\%$	18.20	15.71	18.47	18.57	18.86
$51 \times 51+1\%$	18.03	15.95	18.23	18.93	19.16

Kernel/Noise	Avg. SSIM				
	L0	L0+R	L0+L	L0+P	L0+S
$33 \times 33+3\%$	0.86	0.84	0.87	0.88	0.90
$45 \times 45+2\%$	0.82	0.81	0.83	0.86	0.88
$51 \times 51+1\%$	0.78	0.83	0.79	0.83	0.86

Table 2: The average PSNR and SSIM results of the 20 text images in three kernel-noise settings.

We provide a visual example of kernel/noise setting (iii) in Figure 4, which contains an equation and some English text. As shown by the example, the original L0 method pro-

duces pixel values that do not exist in the original text image, leading to many visible artifacts. A direct rounding to the binary image (L0+R) removes out-of-range pixel values, but it also destroys recovered structures, which lowers both visual quality and quantitative measures. Using the ℓ_0 regularizer in pixel domain (L0+L) slightly improves over the L0 method in PSNR and SSIM, but visual artifacts are still visible in the restored image. On the other hand, regularizers specifically enforcing distinct pixel values including the method of [27] (L0+P) and our regularizer (L0+S) can more effectively incorporate the binary intensity constraint. Our method recovers more structures and suppresses more intensity artifacts with the soft-rounding operator, leading to the best visual quality and quantitative results. More examples can be found in the supplemental material.

4.2. Pattern Image Restoration

Pattern images are an important class of images which can be frequently found from human-made objects, artistic designs and paintings. An important property of pattern images is that their visual contents are represented with a few distinct pixel values [16]. We test our method to remove noise from pattern images. In our experiments, we collect 10 pattern image examples and degrade them with three high Gaussian noise levels (15%, 20% and 25%), such that the restoration performance depends heavily on the regularizers. We adopt the BM3D denoiser as the baseline method [6] and combine it with the soft-rounding regularizer (denoted as BM3D+S). The numerical solution of




Figure 5: *Pattern Image Denoising Examples*. Top row: this pattern image is degraded with 25% Gaussian noise. The distinct pixel values are 32, 76, 142, 230 respectively. Bottom row: this pattern image is degraded with 20% Gaussian noise. The distinct pixel values are 32, 70, 135, 158, 231 respectively. Details are best viewed on screen.

Noise	Avg. PSNR (in dB)		
	BM3D	BM3D+R	BM3D+S
15%	29.77	29.75	31.13
20%	26.71	27.21	28.67
25%	24.59	25.21	25.96

Noise	Avg. SSIM		
	BM3D	BM3D+R	BM3D+S
15%	0.92	0.91	0.94
20%	0.89	0.89	0.92
25%	0.86	0.85	0.88

Table 3: *The average PSNR and SSIM results of the 10 pattern images in three noise settings.*

BM3D+S can be found in the supplemental material. For comparison, we include a post-processing method which directly rounds the BM3D denoised image pixels to nearest known pixel values (denoted as BM3D+R). Note that previous binary intensity regularizers [18, 27] can not be easily extended to handle more than two distinct pixel values.

Table 3 reports the average PSNR and SSIM results over 10 pattern images for the three noise levels. Compared to BM3D, our method (BM3D+S) improves the average PSNR by at least 1.3dB and the average SSIM by at least 0.02 for the three cases. The simple rounding method (BM3D+R) sometimes leads to a deterioration either in PSNR or in SSIM. We provide two visual examples with different noise levels are in Figure 5. As these results show,

the original BM3D method tends to introduce out-of-the-range pixel values and blurred boundaries. The post rounding method (BM3D+R) improves the image contrast with known intensity information, but it also introduces visible errors on region boundaries. On the other hand, our method improves both visual quality and quantitative results with better contrast and image structures.

We further present several deconvolution examples of practical pattern images in Figure 6. We use the L0 deconvolution method as the baseline method and combine the distinct pixel value information with the L0+R and L0+S method. As in the case of denoising, our method (L0+S) significantly improves the visual quality and the quantitative results, while the direct rounding method (L0+R) introduces extra artifacts and removes useful structures in the original images. Please refer to the supplemental material for more examples.

4.3. Natural Image Deconvolution

We randomly select 200 natural images from the BSDS500 dataset [15] and test them with the same kernel/noise settings used in Sec. 4.1. For natural image deconvolution, we choose the total variation method (denoted as TV) [20] as our baseline method. We integrate the integer intensity constraint into the TV method using soft-




Figure 6: *More L0 Deconvolution Examples. Top row: the 1D barcode image is degraded with a 65×65 kernel and 1% Gaussian noise. The distinct pixel values for barcode and background pixels are $t_1 = 0, t_2 = 255$ respectively. Bottom row: the license plate image is degraded with a 55×55 kernel and 3% Gaussian noise. The distinct pixel values are 24, 53, 158, 224, 255 respectively. Details are best viewed on screen.*

Kernel/Noise	# of Images with PSNR Difference D			
	> 0	< 0	> 0.01	< -0.01
$33 \times 33+3\%$	129	71	24	0
$45 \times 45+2\%$	126	74	22	0
$51 \times 51+1\%$	132	68	68	1

Table 4: *Quantitative comparison of TV+S and TV on 200 BSDS500 images for the three kernel/noise settings. The PSNR difference is defined as $D = PSNR_{TV+S} - PSNR_{TV}$.*

rounding (denoted as TV+S) and compare it to the post-rounding method (denoted as TV+R).

Our experiments show that compared to TV, TV+S improves the average PSNR by 0.005dB, 0.005dB and 0.025dB for the three cases, respectively, while TV+R slightly degrades the performance by about 0.002dB for each case. Even though the average contribution of the constraint on integer pixel values seems marginal, a detailed performance breakdown can shed more light on its effect. Specifically, we report the PSNR difference between TV+S and TV on the 200 images in Table 4: for all the three kernel/noise settings, TV+S improves TV on most images; and for at least 11% of the images, the improvement obtained with TV+S over TV is at least 0.01dB; on the other hand, it decreases TV performance significantly (> 0.01 dB) only on one image example from all the cases.

Summary. As all the experimental results show, even though the constraint on pixel values is simple, effectively incorporating it in image restoration often improves the visual quality and the quantitative results for a variety of images and various restoration methods. Furthermore, for special classes of natural images that have distinct pixel values, it leads to considerable performance improvement.

5. Conclusion

In this work, we describe a new regularizer to augment current image restoration algorithms when the original images have known distinct pixel values. Our regularizer is in the form of a function and can be efficiently implemented with a *soft-rounding* operation. This regularizer can be readily incorporated into most existing image restoration methods and our experiments on restoration of binary text images, pattern images with multiple pixel values and natural images show that its incorporation leads to considerable performance improvements to the state-of-the-art methods.

There are a few directions we would like to further improve the current work. Currently, our method relies on the knowledge of the distinct pixel values in the original image. Another scenario is when we only know the exact number of distinct pixel values but not the values themselves. One subsequent study is to augment our algorithm so that it can also simultaneously estimate the distinct pixel values and restore the corrupted image. Furthermore, discriminative image restoration methods [21, 22] seem to be an effective alternative to the generative methods, and it is also of interest to combine our regularizer with such methods.

Acknowledgements

We would like to thank the reviewers for their constructive comments. This work is supported by US National Science Foundation Research Grant (CCF-1319800), US National Science Foundation Early Faculty Career Development (CAREER) Award (IIS-0953373), National Science Foundation of China (61271430) and National Science and Technology Major Project of China (2012ZX01039-004).

References

- [1] M. Afonso, J. Bioucas-Dias, and M. Figueiredo. An augmented lagrangian approach to linear inverse problems with compound regularization. In *ICIP*, pages 4169–4172, 2010.
- [2] M. Afonso, J. Bioucas-Dias, and M. Figueiredo. Fast image recovery using variable splitting and constrained optimization. *TIP*, 19(9):2345–2356, 2010.
- [3] M. R. Banham and A. K. Katsaggelos. Digital image restoration. *IEEE Signal Process. Mag.*, 14(2):24–41, 1997.
- [4] X. Chen, X. He, J. Yang, and Q. Wu. An effective document image deblurring algorithm. In *CVPR*, pages 369–376, 2011.
- [5] T. S. Cho, N. Joshi, C. L. Zitnick, S. B. Kang, R. Szeliski, and W. T. Freeman. A content-aware image prior. In *CVPR*, pages 169–176, 2010.
- [6] K. Dabov, A. Foi, V. Katkovich, and K. Egiazarian. Image denoising by sparse 3-d transform-domain collaborative filtering. *TIP*, 16(8):2080–2095, 2007.
- [7] D. L. Donoho. De-noising by soft-thresholding. *TIT*, 41(3):613–627, 1995.
- [8] D. Geman and C. Yang. Nonlinear image recovery with half-quadratic regularization. *TIP*, 4(7):932–946, 1995.
- [9] F. Heide, M. Steinberger, Y.-T. Tsai, M. Rouf, D. Pajak, D. Reddy, O. Gallo, J. Liu, W. Heidrich, K. Egiazarian, J. Kautz, and K. Pulli. FlexISP: A flexible camera image processing framework. In *ACM SIGGRAPH Asia*, pages 231:1–231:13, 2014.
- [10] N. Joshi, C. L. Zitnick, R. Szeliski, and D. J. Kriegman. Image deblurring and denoising using color priors. In *CVPR*, pages 1550–1557, 2009.
- [11] D. Krishnan and R. Fergus. Fast image deconvolution using hyper-Laplacian priors. In *NIPS*, pages 1033–1041, 2009.
- [12] W.-S. Lai, J.-J. Ding, Y.-Y. Lin, and Y.-Y. Chuang. Blur kernel estimation using normalized color-line prior. In *CVPR*, pages 3753–3761, 2015.
- [13] A. Levin, R. Fergus, F. Durand, and W. T. Freeman. Image and depth from a conventional camera with a coded aperture. *TOG*, 26(3):70:1–70:9, 2007.
- [14] A. Levin, Y. Weiss, F. Durand, and W. T. Freeman. Understanding and evaluating blind deconvolution algorithms. In *CVPR*, pages 1964–1971, 2009.
- [15] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *ICCV*, pages 416–423, 2001.
- [16] X. Mei, W. Dong, B.-G. Hu, and S. Lyu. UniHIST: a unified framework for image restoration with marginal histogram constraints. In *CVPR*, pages 3753–3761, 2015.
- [17] J. Nocedal and S. Wright. *Numerical Optimization*, 2nd edition. Springer, 2006.
- [18] J. Pan, Z. Hu, Z. Su, and M.-H. Yang. Deblurring text images via L0-regularized intensity and gradient prior. In *CVPR*, pages 2901–2908, 2014.
- [19] N. Parikh and S. Boyd. Proximal algorithms. *Foundations and Trends in Optimization*, 1(3):123–231, 2013.
- [20] L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1):259–268, 1992.
- [21] U. Schmidt and S. Roth. Shrinkage fields for effective image restoration. In *CVPR*, pages 2774–2781, 2014.
- [22] U. Schmidt, C. Rother, S. Nowozin, J. Jancsary, and S. Roth. Discriminative non-blind deblurring. In *CVPR*, pages 604–611, 2013.
- [23] M. J. Wainwright and E. Simoncelli. Scale mixtures of Gaussians and the statistics of natural images. In *NIPS*, pages 855–861, 2000.
- [24] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli. Image quality assessment: from error visibility to structural similarity. *TIP*, 13(4):600–612, 2004.
- [25] L. Xu, C. Lu, Y. Xu, and J. Jia. Image smoothing via L0 gradient minimization. In *ACM SIGGRAPH Asia*, pages 174:1–174:12, 2011.
- [26] L. Xu, S. Zheng, and J. Jia. Unnatural l0 sparse representation for natural image deblurring. In *CVPR*, pages 1107–1114, 2013.
- [27] J. Zhang. An alternating minimization algorithm for binary image restoration. *TIP*, 21(2):883–888, 2012.
- [28] D. Zoran and Y. Weiss. From learning models of natural image patches to whole image restoration. In *ICCV*, pages 479–486, 2011.
- [29] W. Zuo, L. Zhang, C. Song, D. Zhang, and H. Gao. Gradient histogram estimation and preservation for texture enhanced image denoising. *TIP*, 23(6):2459–2472, 2014.