A Variational Approach for Exact Histogram Specification*

Raymond Chan[1], Mila Nikolova[2], and You-Wei Wen[1,3]

[1] Department of Mathematics, The Chinese University of Hong Kong, Shatin, NT, Hong Kong.
[2] Centre de Mathématiques et de Leurs Applications, ENS de Cachan, 61 av. du Président Wilson, 94235 Cachan Cedex, France.
[3] Department of Mathematics, Kunming University of Science and Technology, Yunnan, P.R.China.

Abstract. We focus on exact histogram specication when the input image is quantified. The goal is to transform this input image into an output image whose histogram is exactly the same as a prescribed one. In order to match the prescribed histogram, pixels with the same intensity level in the input image will have to be assigned to different intensity levels in the output image. A novel method enabling to order strictly all pixel values was proposed to solve the problem by using auxiliary attributes to classify pixels with the same intensity value. Local average intensities and wavelet coefficients have been used by the past as the second attribute. However, these methods cannot enable strict-ordering without degrading the image. In this paper, we propose a variational approach to establish an image preserving strict-ordering of the pixel values. We show that strict-ordering is achieved with probability one. Our method is image preserving in the sense that it reduces the quantization noise in the input quantified image. Numerical results show that our method gives better quality images than the preexisting methods.

Key words: Exact histogram specification, strict-ordering, variational methods, restoration from quantization noise, smooth nonlinear optimization, convex minimization

1 Introduction

Image histogram processing is the act of altering each individual pixel of an image by modifying its dynamic range in order to improve the contrast of the whole image. It is an important image processing task with many real-world applications, such as contrast enhancement, segmentation, watermarking, among many others.

* The research was supported in part by HKRGC Grant CUHK400510 and DAG Grant 2060408. E-mail: rchan@math.cuhk.edu.hk,nikolova@cmla.ens-cachan.fr, wenyouwei@gmail.com.

In histogram processing, image intensity level is viewed as a random variable characterized by its probability density function. The histogram of an image shows the empirical distribution of the intensity levels of its pixels. One of the basic histogram processing problem is histogram equalization [10, 18]. It aims to find a transformation so that the output image has a uniform histogram. In the continuous setting the random variable defined by the cumulative distribution function of the intensity levels is uniformly distribution in $[0, 1]$, and hence such a function can always be found. More generally, we may want to yield an output image with pre-specified histogram shapes. This problem is called *histogram specification* or *histogram matching*. The prescribed histogram can be given according to various needs. For example, it can be the histogram of another image, a modified version of the original histogram [19], or a "weighted" histogram of two histograms [6, 7].

Numerous methods have been proposed to modify the histogram of an input image. The simplest method is histogram linear stretching [13]. Histogram clipping method [19] limits the maximum number of pixels for each intensity level to a given constant and the clipped pixels are then uniformly distributed among the other intensity levels with pixels less than the clip limit. Several other methods were proposed to preserve the mean brightness of the input image [3, 12, 23]. In [20], Sapiro and Caselles proposed histogram modification via image evolution equations. Arici *et al.* proposed a general framework for histogram modification [1].

The principle behind histogram specification methods is straightforward for real-valued (analog) images: the histogram of the input image and the prescribed histogram should be equalized to uniform distribution first, say by T_i and T_t respectively. Then the output image can be obtained from the composite transformation $T_t^{-1} \circ T_i$. Since the images are real-valued, T_i and T_t are one-to-one functions, and hence $T_t^{-1} \circ T_i$ is well-defined. The principle fails, however, for quantized (digital) images, which is the case of all digital video systems. The reason is that for quantized images, the intensity levels of all pixels take a limited number of discrete values. Therefore their cumulative density functions are staircase functions rather than strictly increasing functions like those for the real-valued images. Indeed, there are groups of pixels with the same intensity value. Some pixels in such a group will have to be mapped to pixels with different intensity values to match the prescribed histogram. This task cannot be achieved without the use of some auxiliary information on pixel values.

Methods to obtain strict ordering based on a quantized image were proposed in [4, 5, 22]. Once all pixels are ordered strictly, the prescribed intensity values are assigned exactly according to histogram specification.

Then the prescribed intensity value is assigned according to the resulting ordering. Coltuc *et al.* considered to use the average intensities of neighboring pixels as the auxiliary attribute [5]. Considering two pixels with the same intensity value, the mean values over the neighborhoods centered on each pixel are compared to order these two pixels. If the mean values are still the same, then they choose larger neighborhoods and continue in the same way until all pixels are ordered. Wan and Shi argued that the local mean approach fails to sharpen the edges of the output image [22]. They proposed to order the pixels according to the absolute values of its wavelet coefficients. The wavelet-based approach tends to amplify the noise since a noise in a smooth region may be mistaken as an edge and hence is sharpened. Post-processing approach or iterative methods can be applied to suppress the amplified noises [2]. We emphasize that both the local mean approach and the wavelet-based approach cannot realize strict ordering without degrading the input quantized image. This is a major drawback.

In this paper, we propose a variational method that enables to order strictly the pixel values of a quantified image by restoring it from the quantization noise. We prove that the pixels of the restored image can be totally-ordered with probability equal to one. Our experimental results show that the proposed method is very efficient and produces images of better quality than both the local mean method [5] and the wavelet-based method [22].

The outline of the paper is as follows. In Section 2, we present the proposed method. In Section 3, numerical examples are given to demonstrate the effectiveness of the proposed model. Concluding remarks are given in Section 4.

2 Variational Approach for Exact Histogram Specification

In this section, we introduce the definition of strict-ordering and then we propose our variational approach for exact histogram specification. First, let us present the problem of exact histogram specification.

Consider an M-by-N input quantized image u whose pixel values live in the set $\mathcal{P} = \{p_1, \cdots, p_L\}$. We assume, without loss of generality, that p_i are in increasing order. For 8-bit images, $\mathcal{P} = \{0, \cdots, 255\}$. Let the grid of u be denoted by

$$\Omega := \{\mathbf{x} : \mathbf{x} = (i, j),\ 1 \le i \le M, 1 \le j \le N\}.$$

The intensity of u at the pixel $\mathbf{x}$ is given by $u_{\mathbf{x}}$. Define

$$\Omega_k := \{\mathbf{x} \in \Omega : u_{\mathbf{x}} = p_k\}, \quad k = 1, 2, \cdots, L.$$

The associated histogram of u is the L-tuple $(|\Omega_1|, |\Omega_2|, \ldots, |\Omega_L|)$, where $|\cdot|$ denotes the cardinality of the set. The problem of exact histogram specification that we consider can be stated as follows: given the input image u, obtained from an original real-valued (analog) image u_{o} by quantization, and a pre-specified histogram $\boldsymbol{h} = (h_1, h_2, \ldots, h_L)$, find an output image v such that its histogram is $\boldsymbol{h}$ and for any $\mathbf{x}, \mathbf{y} \in \Omega$, we have $v_{\mathbf{x}} \le v_{\mathbf{y}}$ if $u_{\mathrm{o},\mathbf{x}} \le u_{\mathrm{o},\mathbf{y}}$.

2.1 Sorting Algorithms

Since $MN \gg L$ generally, there are many pixels that share the same intensity value. In order to order strictly the pixels with the same intensity, auxiliary information must be used. Combining the auxiliary information, we can create a K-vector defined as $\left(u_{\mathbf{x}}, \kappa_{\mathbf{x}}^1, \ldots, \kappa_{\mathbf{x}}^{K-1}\right)$ for $\mathbf{x} \in \Omega$, where $\kappa_{\mathbf{x}}^i \in \mathbb{R}$ is the i-th auxiliary information of the pixel $\mathbf{x}$. Our approach to determine the auxiliary information will be outlined later.

Now we can define an ascending ordering "$\prec$" for pixels in Ω based on such K-tuples. To facilitate the discussions, let $\kappa_{\mathbf{x}}^0 := u_{\mathbf{x}}$. For any two pixels $\mathbf{x}$ and $\mathbf{y}$ in Ω, we say that $\mathbf{x} \prec \mathbf{y}$ if for some $0 \leq \ell \leq K-1$

$$\kappa_{\mathbf{x}}^j = \kappa_{\mathbf{y}}^j \text{ for all } \ 0 \leq j \leq \ell - 1 \text{ and } \kappa_{\mathbf{x}}^\ell < \kappa_y^\ell. \tag{1}$$

For good choices of auxiliary information and K sufficiently large, one can in principle sort all pixels $\mathbf{x}$ in Ω according to the ordering $\prec$. That is, we can order the pixels $\mathbf{x}$ in Ω in such a way that $\mathbf{x}_1 \prec \mathbf{x}_2 \prec \ldots \prec \mathbf{x}_{NM}$.

Once such a strict-ordering is obtained, matching the input histogram to the prescribed one is straightforward. This can be done by dividing the ordered list from left to right into L groups. Starting from $\mathbf{x}_1$ on the list, the first h_1 pixels belong to the first group, and are assigned the intensity of p_1. The next h_2 pixels belong to the second group and are assigned the intensity of p_2, and so on until all pixels are assigned to their new intensities.

Several ideas have been proposed for the auxiliary information. Coltuc *et al.* proposed to use the local average intensities of a pixel's neighborhood as auxiliary information [5]. For pixels having the same intensity, if the average intensities of their neighborhoods are the same, then a larger neighborhood will be chosen to compute the average intensity. This procedure is repeated until all pixels are ordered. The author claimed that $K = 6$ is appropriate for any application. Wan and Shi proposed to order the pixels according to the absolute values of the wavelet coefficients of the whole image [22]. Here we propose a variational approach to obtain pertinent auxiliary information.

2.2 A Variational Approach

Let us emphasize that the input (digital) image u is obtained from an original real-valued (analog) image u_{o} by quantization. Since the pixels of u_{o} have a continuous range, they can be totally-ordered with probability one. The input image u contains quantization noise. The most natural way to define the ordering for the pixels of u is to restore the original real-valued image u_{o} using u and a good prior knowledge. Such a restoration can efficiently be done using a detail preserving variational method as the one we are proposing here.

For any $\mathbf{x} \in \Omega$, let $\mathcal{N}_{\mathbf{x}} \subset \Omega$ be the set of neighboring pixels of $\mathbf{x}$ (in our experiment, we choose $\mathcal{N}_{\mathbf{x}}$ to be the four neighboring pixels of $\mathbf{x}$ in the vertical and horizontal directions). Now we order the pixels by minimizing f in the cost functional $\mathcal{J} : \mathbb{R}^{M \times N} \times \mathbb{R}^{M \times N} \to \mathbb{R}$ given below

$$\mathcal{J}(f, u) = \sum_{\mathbf{x} \in \Omega} \left(\psi\big(f(\mathbf{x}) - u(\mathbf{x})\big) + \beta \sum_{\mathbf{y} \in \mathcal{N}_{\mathbf{x}}} \phi\big(f(\mathbf{x}) - f(\mathbf{y})\big) \right). \tag{2}$$

Here $\beta > 0$ is the regularization parameter and

H1 *$\phi : \mathbb{R} \mapsto \mathbb{R}$ and $\psi : \mathbb{R} \mapsto \mathbb{R}$ are even functions in $\mathcal{C}^s$ with $s \geq 2$, such that $\phi''(t) > 0$ and $\psi''(t) > 0$, $\forall t \in \mathbb{R}$.*

For instance we can choose

$$\psi(t) = \sqrt{t^2 + \alpha_1} \quad \text{and} \quad \phi(t) = \sqrt{t^2 + \alpha_2}\ , \quad \alpha_1 > 0,\ \alpha_2 > 0 \tag{3}$$

which are $\mathcal{C}^\infty$ and analytic. The minimizer of $\mathcal{J}$ in (2) is denoted by $\hat{f}$.

We know that the quantization noise is bounded, $\|u_\text{o} - u\|_\infty \leq 0.5$. This constraint should not be used explicitly however because many pixels may then be stuck on the box constraint which will make strict ordering impossible. Instead, the constraint can be satisfied in a relaxed way by using a slightly smoothed ℓ_1 data-fidelity term like ψ in (3) for $\alpha_1 \gtrapprox 0$ and $\beta \gtrapprox 0$ in (2). By choosing $\beta \gtrapprox 0$, data-fidelity is enhanced. If $\psi(t) = |t|$, some data entries would be kept intact [15] and since data-fidelity is enhanced we would find $\hat{f} = u$. But taking ψ as in (3) for $\alpha_1 \gtrapprox 0$ entails that $\hat{f} \overset{\approx}{\neq} u$. A prior holding for large classes of natural images is that they are almost nowhere constant (see [11]) and that they involve edges and fine structures. Nowhere constant implies that ϕ must be smooth at the origin [16]. For edges and fine structures, ϕ must be affine or nonconvex away from the origin. Since pixels must change no more than $|0.5|$ for an image range equal to 255, the best choice is a convex ϕ of the form (3) for $\alpha_2 \gtrapprox 0$. Below we show that the pixels of $\hat{f}$ can be ordered with probability one.

Definition 1. *A function $\mathcal{F} : O \mapsto \mathbb{R}^{M \times N}$, where O is an open domain in $\mathbb{R}^{M \times N}$, is said to be a* minimizer function *relevant to $\mathcal{J}(\cdot, O)$ if for every $u \in O$, the point $\hat{f} = \mathcal{F}(u)$ is a strict local minimizer of $\mathcal{J}(\cdot, u)$.*

For any $u \in \mathbb{R}^{M \times N}$, the functional $\mathcal{J}(\cdot, u)$ in (2), satisfying H1, is strictly convex and coercive, hence for any u and $\beta > 0$, it has a unique minimizer. What is more, one can show that $\mathcal{J}$ has a unique minimizer function $\mathcal{F} : \mathbb{R}^{M \times N} \mapsto \mathbb{R}^{M \times N}$ which is $\mathcal{C}^{s-1}$ continuous, see [14].

We denote by $\mathbb{L}^{M \times N}$ the Lebesgue measure on $M \times N$ subsets of matrices using the isomorphism between $M \times N$ real matrices and MN-length real vectors. Our main theoretical results , proven in [14], are summarized below. The components of the minimizer function $\mathcal{F}$ are denoted by $\mathcal{F}_\mathbf{x}$, $\mathbf{x} \in \Omega$.

Theorem 1. *Let $\mathcal{J}$ in (2) satisfy H1. For its minimizer function $\mathcal{F} : \mathbb{R}^{M \times N} \mapsto \mathbb{R}^{M \times N}$, define the sets $\mathcal{Q}$ and $\mathcal{R}$ as follows:*

$$\mathcal{Q} = \{u \in \mathbb{R}^{M \times N}\ :\ \mathcal{F}_\mathbf{x}(u) = \mathcal{F}_\mathbf{y}(u)\ , (\mathbf{x}, \mathbf{y}) \in \Omega \times \Omega,\ \mathbf{x} \neq \mathbf{y}\}\ , \tag{4}$$

$$\mathcal{R} = \{u \in \mathbb{R}^{M \times N}\ :\ \mathcal{F}_\mathbf{x}(u) = u_\mathbf{y}\ , (\mathbf{x}, \mathbf{y}) \in \Omega \times \Omega,\ \mathbf{x} \neq \mathbf{y}\}\ . \tag{5}$$

The sets $\mathcal{Q}$ and $\mathcal{R}$ are closed, and satisfy $\mathbb{L}^{M \times N}(\mathcal{Q}) = 0$ and $\mathbb{L}^{M \times N}(\mathcal{R}) = 0$.

The set $\mathcal{Q}$ in (4) contains all possible $u \in \mathbb{R}^{M \times N}$ such that the minimizer $\hat{f} = \mathcal{F}(u)$ might have two equal entries, $\mathcal{F}_\mathbf{x}(u) = \mathcal{F}_\mathbf{y}(u)$ for some $\mathbf{x} \neq \mathbf{y}$ belonging to Ω. The set $\mathcal{R}$ in (5) contains all possible $u \in \mathbb{R}^p$ such that the minimizer $\hat{f} = \mathcal{F}(u)$ might contain some quantized entries, $\mathcal{F}_{\mathbf{x}}(u) = u_{\mathbf{y}}$ for some $\mathbf{x}, \mathbf{y} \in \Omega$.

Even though $\mathcal{Q}$ is not empty, since $\mathcal{Q}$ is closed and of null Lebesgue measure, the chance that real-world quantized images u live in it, is *null*. Thus, $\mathcal{F}_{\mathbf{x}}(u) \neq \mathcal{F}_{\mathbf{y}}(u)$, for $\mathbf{x} \neq \mathbf{y}$, is a *generic* property of the minimizers $\mathcal{F}$ of $\mathcal{J}$, as given in (2) and satisfying H1. *For any real-world quantized image u, the entries of the minimizer $\hat{f} = \mathcal{F}(u)$ can be classified with probability one.* In the numerous experiments we have done, we never found natural quantized images belonging to $\mathcal{Q}$ nor to $\mathcal{R}$, i.e. in all cases we could perfectly order the pixels of $\hat{f}$.

There are many methods to compute the minimizer $\hat{f}$ of $\mathcal{J}(\cdot, u)$ in (2) [8, 9, 17, 21]. We applied fixed point iteration method [21] to find $\hat{f}$. Once we have find the minimizer $\hat{f}$, we establish the ordering of the pixels based on the 2-tuple $(u_{\mathbf{x}}, \hat{f}_{\mathbf{x}})$ to produce the quantized output image v.

3 Experimental Results

The performance of the proposed method for exact histogram specification was evaluated using extended numerical experiments. Some of them are presented below. We compare our method with the local mean (LM) algorithm [5] for $K = 6$ as recommended by the authors and with the wavelet-based algorithm (WA) in [22]. For our method, we set α_i, $i = 1, 2$, in (3) to 0.01, and β in (2) to 0.1. We stop the iteration when the relative difference between the iterant is less than 10^{-8}.

In order to measure the results quantitatively, we start out with a given true quantized image w with histogram $\boldsymbol{h}_w$; then we degrade it to obtain an input quantized image u. By applying the three methods on u with prescribed histogram $\boldsymbol{h}_w$, we obtain an output image v which is in fact a restored version of w. We use peak-signal-to-noise-ratio to measure the quality of the output image v with respect to w. It is defined as $\text{PSNR} = 20 \log_{10}(255NM/\|v - w\|_2)$. We tried two sets of degradation to obtain the input image u.

3.1 Contrast Compression

In our first set of degradation, the true quantized images w are chosen to be the 256-by-256 8-bit images of "Cameraman", "Lenna" and "Peppers". The input image u is obtained from w by the degradation: $u = \text{round}(\rho \cdot w)$, where $\rho < 1$ is a constant. This situation arises when a picture is taken with insufficient exposure time, or when we want to compress the image by reducing the number of intensity levels. For example, a 7-bit image can be obtained from an 8-bit image by using $\rho = 0.5$. The input images u for $\rho = 0.3$ are shown in the first row of Figure 1. In the tests, we used LM, WA and our method to obtain the output images v having a prescribed histogram $\boldsymbol{h}_w$.

The comparisons of LM, WA and our algorithm are shown in Table 1. We see from the PSNR values that our method outperforms LM and WA in all cases. In order to save space, we just show the output images v by our method, see

Fig. 1. First row: the input images. Second row: the output images by our method.

the second row of Figure 1. The difference images between the true image w and the output image v are shown in Figure 2. We can discern more features in the first row and the second row than in the third row. It demonstrates that our algorithm yields the best restoration.

	Cameraman			Lenna			Peppers		
ρ	LM	WA	Ours	LM	WA	Ours	LM	WA	Ours
0.8	55.97	55.86	56.07	55.64	55.50	55.73	55.98	55.68	56.05
0.7	54.15	54.07	54.33	53.93	53.77	53.98	54.17	53.82	54.24
0.6	52.96	52.84	53.09	52.69	52.50	52.74	52.93	52.64	53.02
0.5	51.93	51.84	52.06	51.67	51.51	51.74	51.97	51.66	52.04
0.4	49.24	49.12	49.45	49.01	48.72	49.15	49.40	48.90	49.58
0.3	47.16	46.98	47.42	46.74	46.32	47.00	47.19	47.42	47.50
0.2	44.07	33.87	44.46	43.70	43.14	44.07	44.39	43.57	44.94
0.1	38.75	38.55	39.36	38.72	37.83	39.38	39.44	38.36	40.38

Table 1. The PSNR (dB) between the true image w and the output image v.

One important indicator for a good exact histogram specification algorithm is to see if it can establish a strict ordering for all the pixels. If a sorting method yields two pixels sharing the same value, we consider this as a failure of the method. We call a pair-pixel for any two pixels having the same value at the issue of the sorting algorithm. Table 2 shows the numbers of pair-pixels produced by the three methods. We find that LM and WA have a high number of pair-pixels while our method can give a total ordering of all pixels for all three images. Incidentally, for the "Cameraman" image, there are 13,859 pair-pixels for $\rho = 0.1$

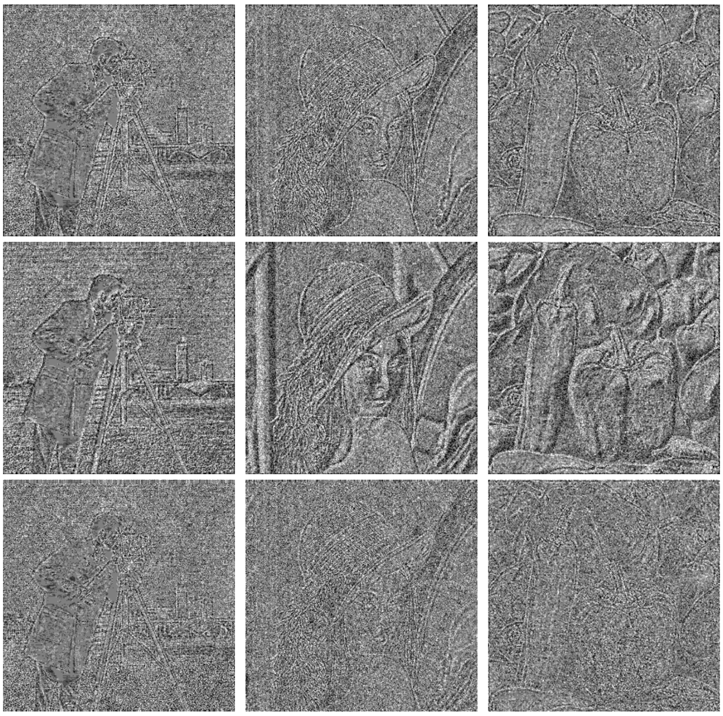

Fig. 2. The difference images between the true quantized image w and the output image v. First row: LM method. Second row: WA method. Third row: our method.

for WA. Compared with the image size, which has 65,532 pixels, the ordering failure rate is about 21%.

3.2 Histogram Equalization Inversion

The second set of degradation is done as follows. Given the true quantized image w with histogram $\boldsymbol{h}_w$, we apply each individual method to get the pixel ordering of w. Then we use the ordering to match w to an image with uniform histogram. The resulting image is used as the input image u of our experiment. Given u and the prescribed histogram $\boldsymbol{h}_w$, we apply each individual method to obtain the output image v. If the ordering among the pixels is preserved by the method, we should have $v = w$ exactly.

For this experiment, we tried the three images in Section 3.1 together with 15 real 768-by-512 8-bit images available at http://r0k.us/graphics/kodak/. Color images are converted to the gray-scale images first. Table 3 shows the PSNR of the results by the three methods. Figures 3–5 give the difference images between w and v on "Cameraman", "Lenna", "Peppers" and two of the 15 images. We notice from Table 3 that WA method yields better PSNR than LM method in all images, but worse than our method in all cases except for the "Lenna" and

	Cameraman			Lenna			Peppers		
ρ	LM	WA	Ours	LM	WA	Ours	LM	WA	Ours
0.8	3	154	0	0	7	0	0	13	0
0.7	90	377	0	0	8	0	0	15	0
0.6	88	437	0	0	15	0	0	35	0
0.5	76	587	0	0	26	0	1	38	0
0.4	344	1,267	0	1	66	0	1	145	0
0.3	829	2,293	0	1	177	0	20	403	0
0.2	2,146	4,529	0	36	803	0	109	1,205	0
0.1	6,517	13,859	0	1,493	5,499	0	3,211	7,230	0

Table 2. The numbers of pair-pixels from the three methods.

"Peppers" images. Though WA method yields better PSNR than our method in those two cases, from the figures, we can discern more features in the difference images by WA method than by our method. This indicates that our method is more accurate.

Image	LM	WA	Ours	Image	LM	WA	Ours
Cameraman	48.25	48.44	48.79	Kadim07	43.74	43.83	48.09
Lenna	51.24	51.75	51.50	Kadim08	48.33	48.55	50.77
Peppers	51.99	52.66	52.14	Kadim09	44.85	44.94	48.71
Kadim01	41.77	41.81	43.36	Kadim10	44.74	44.85	47.29
Kadim02	43.32	43.38	45.12	Kadim11	45.26	45.35	46.63
Kadim03	44.69	44.76	47.95	Kadim12	40.66	40.70	45.64
Kadim04	45.92	45.99	46.86	Kadim13	47.42	47.58	50.39
Kadim05	49.41	49.71	49.81	Kadim14	45.76	45.86	47.19
Kadim06	44.88	44.95	48.80	Kadim15	49.00	49.23	49.71

Table 3. The PSNR (dB) between the true image w and output images v.

4 Conclusions

In this paper, we propose a variational approach for exact histogram specification. Since the energy we minimize is smooth, its minimizers enable to order strictly all the pixels in the image. Noticing also that our method reduces the quantification noise, the obtained results outperform the preexisting methods.

References

1. T. Arici, S. Dikbas, and Y. Altunbasak. A histogram modification framework and its application for image contrast enhancement. *IEEE Transactions on Image Processing*, 18(9):1921–1935, 2009.

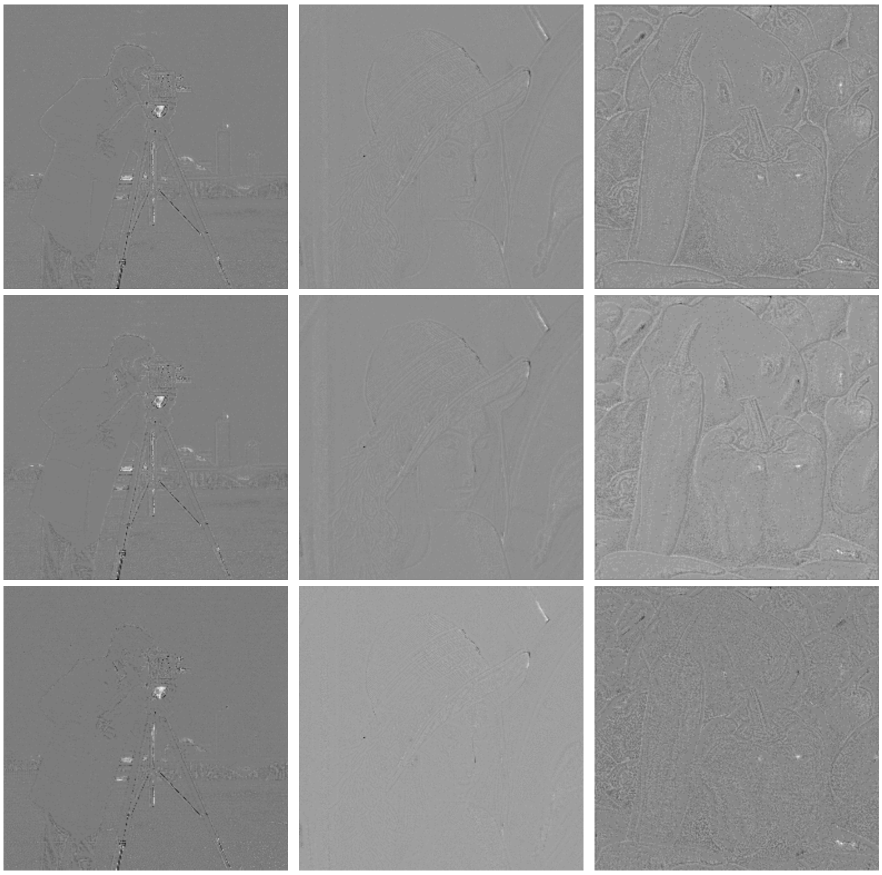

Fig. 3. The difference image between w and v by LM (first row), WA (second row) and our method (third row). Our method yields fewest features in the difference images.

2. A.N. Avanaki. Exact global histogram specification optimized for structural similarity. *Optical review*, 16(6):613–621, 2009.
3. S. Chen and A. R. Ramli. Contrast enhancement using recursive mean-separate histogram equalization for scalable brightness preservation. *IEEE Transactions on Consumer Electronics*, 49(4):1301–1309, 2003.
4. D. Coltuc and P. Bolon. An inverse problem: Histogram equalization. In *Signal Processing IX, Theories and Applications, EUSIPCO'98*, volume 2, pages 861–4, 1998.
5. D. Coltuc, P. Bolon, and J.-M. Chassery. Exact histogram specification. *IEEE Transactions on Image Processing*, 15(5):1143–1152, 2006.
6. I.J. Cox, S. Roy, and S.L. Hingorani. Dynamic histogram warping of image pairs for constant image brightness. In *Image Processing, 1995. Proceedings., International Conference on*, volume 2, pages 366–369. IEEE, 2002.
7. J. Delon. Midway image equalization. *Journal of Mathematical Imaging and Vision*, 21(2):119–134, 2004.

Fig. 4. Top-left corner: the given true quantized image w. The difference image between w and the output image v by LM method (top-right), WA method (bottom-left) and our method (bottom-right). Our method yields fewest features in the difference images.

8. D. Geman and G. Reynolds. Constrained restoration and the recovery of discontinuities. *IEEE Trans. Pattern Anal. Mach. Intell.*, 14(3):367–383, 1992.
9. D. Geman and C. Yang. Nonlinear image recovery with half-quadratic regularization. *IEEE Trans. Image Process.*, 4(7):932–946, JUL 1995.
10. R.C. Gonzalez and R.E. Woods. Digital image processing. *Beijlng: Publishing House of Electronics Industry*, 2005.
11. Y. Gousseau and J.-M. Morel, *Are natural images of bounded variation?*, SIAM Journal on Mathematical Analysis, (2001).
12. Y. Kim. Contrast enhancement using brightness preserving bi-histogram equalization. *IEEE Transactions on Consumer Electronics*, 43(1):1–8, 1997.
13. S. Narasimhan and S. Nayar. Contrast restoration of weather degraded images. *IEEE Trans. Pattern Anal. Mach. Intell.*, 25:713–724, 2003.
14. M. Nikolova. On the minimizers of smooth detail preserving regularized objectives. *Report 2011.*
15. M. Nikolova, *A variational approach to remove outliers and impulse noise*, *Journal of Mathematical Imaging and Vision*, 20 (2004).
16. M. Nikolova, *Weakly constrained minimization. Application to the estimation of images and signals involving constant regions*, *Journal of Mathematical Imaging and Vision*, 21 (2004), pp. 155–175.
17. M. Nikolova and M. Ng. Analysis of half-quadratic minimization methods for signal and image recovery. *SIAM J. Sci. Comput.*, 2005.
18. W. K. Pratt. *Digital Image Processing.* 2nd ed., Wiley, New York, 1991.

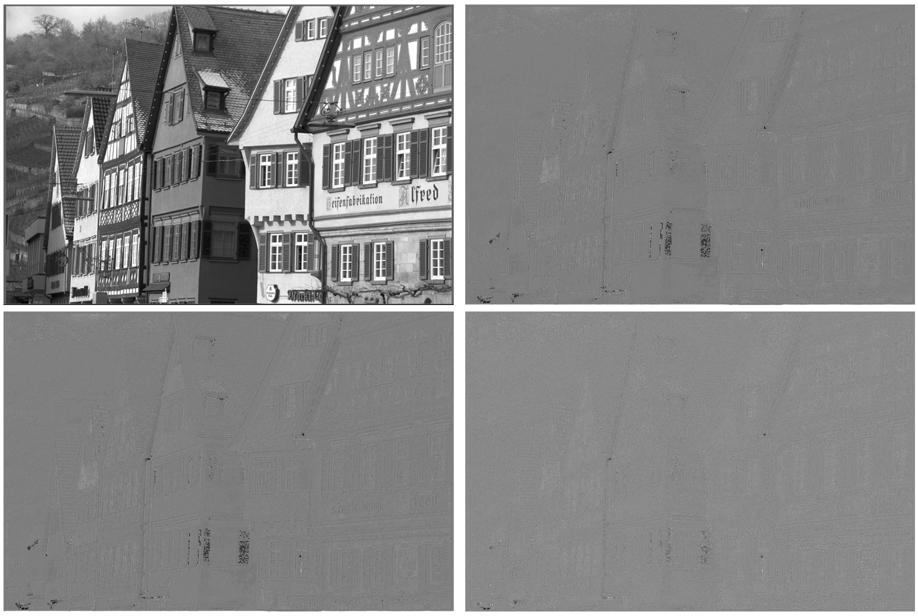

Fig. 5. Top-left corner: the given true image w. The difference image between w and the output image v by LM method (top-right), WA method (bottom-left) and our method (bottom-right). Our method yields fewest features in the difference images.

19. A.M. Reza. Realization of the contrast limited adaptive histogram equalization (CLAHE) for real-time image enhancement. *The Journal of VLSI Signal Processing*, 38(1):35–44, 2004.
20. G. Sapiro and V. Caselles. Histogram modification via partial differential equations. In *Image Processing, 1995. Proceedings., International Conference on*, volume 3, pages 632–635. IEEE, 2002.
21. C. Vogel and M. Oman. Iterative method for total variation denoising. *SIAM J. Sci. Comput.*, 17:227–238, 1996.
22. Y. Wan and D. Shi. Joint exact histogram specification and image enhancement through the wavelet transform. *IEEE Transactions on Image Processing*, 16(9):2245–2250, 2007.
23. C. Wang and Z. Ye. Brightness preserving histogram equalization with maximum entropy: a variational perspective. *IEEE Transactions on Consumer Electronics*, 51(4):1326–1334, 2005.