COLOR DE-QUANTIZING THROUGH ITERATED DYNAMIC HARD THRESHOLDING ON AN OVERCOMPLETE REPRESENTATION

Luis Mancera[†], *Javier Portilla*[♯]

[†]Dept. Computer Science and AI
Universidad de Granada
18071 Granada, Spain
mancera@decsai.ugr.es

[♯]Instituto de Optica
CSIC
28006, Madrid, Spain
portilla@io.cfmac.csic.es

ABSTRACT

We propose a new color de-quantizing method for paletted images based on maximizing the sparseness of the overcomplete wavelet analysis of the estimation within the consistency set defined by the observation. The sparsity is enforced by minimizing the ℓ_0-norm of the coefficients of the wavelet analysis representation. The resulting method iterates between hard-thresholding of the linear response of current estimate and a new proposed consistency projection onto the Voronoi cells defined by color quantization. Results indicate that our method outperforms those based on linear diffusion and minimization of ℓ_1-norm.

1. INTRODUCTION

Color quantization reduces the number of colors in a RGB image by replacing each of them with that color from a given representative set (palette) which minimizes some distance measure. We call de-quantizing to the process of estimating the original image from a quantized version. Usually, the artifacts derived from quantization are close or even below the visibility threshold, but they can become evident in a number of situations. For example, if the local luminance range is stretched for detail inspection, or as a previous step for blurred image restoration when the main source of noise is quantization. It can be also used previously to feature extraction, like the gradient of the luminance.

De-quantizing in a transformed domain has been widely approached, specially in the context of post-processing compressed images (e.g., [2, 3]). Only recently there has been a growing interest in approaching the problem in the image domain. Up to our knowledge, current existing techniques for color de-quantizing are based in constrained diffusion (e.g., [4, 5]). However, this type of strategies, although resulting in efficient algorithms, are too simple to provide satisfactory results. It is known that natural images typically produce sparse distributions of their wavelet coefficients, i.e., their energy is mostly concentrated in a small proportion of coefficients [6]. This knowledge has been often incorporated to the image prior by using Generalized Gaussian distributions (GGD) with exponent p to characterize the wavelet response (e.g. [7, 8]). Within GGDs, convex versions ($p \geq 1$) have been traditionally more popular because global minimum can be achieved efficiently. However, local solutions with non-convex GGD priors ($0 \leq p < 1$) have been recently shown to provide better performance in a wide range of image processing applications (e.g., [9, 10]).

This paper proposes a new color image de-quantizing method based on maximizing the sparseness in the overcomplete wavelet linear representation of images consistent with the observation. The sparsity is enforced by using a GGD prior with $p = 0$ for the coefficients of the wavelet analysis representation. The resulting method iterates between hard-thresholding of the linear response of the current estimate and a new proposed consistency projection onto the Voronoi cells defined by color quantization. Through real examples, results indicate that our proposed method outperforms those based on linear diffusion strategies and also methods based on minimization of ℓ_1-norm.

2. COLOR DEQUANTIZING BY SPARSE AND CONSISTENT APPROXIMATION

Let $\mathbf{X}$ be an original RGB image of size $N \times 3$, where each column is a lexicographically ordered color channel. We denote $\mathbf{X}_i$ to the color vector at the i-th pixel of $\mathbf{X}$. We denote $\mathbf{P}$ to the pallete, which is a $C \times 3$ matrix where C representative colors are stored. Given some distance measure, $d(a, b)$, between two colors a and b (e.g., Euclidean distance), then the color space is partitioned into C non-overlapped Voronoi regions $\{V_{P_k}\}_{k=1,\cdots,C}$, each one defined as:

$$V_{P_k} = \{\mathbf{X}_i \in R^3 : d(\mathbf{X}_i, \mathbf{P}_k) \leq d(\mathbf{X}_i, \mathbf{P}_m)\ \forall\, m = 1, \ldots, C\}$$

Then, the observed image $\mathbf{Y}$ is obtained by $\mathbf{Y} = f(\mathbf{X})$, where each $f(\mathbf{X}_i) = \mathbf{P}_k$ if $\mathbf{X}_i \in V_{P_k}$. Note that each color

This work has been supported by the Spanish research programme Consolider Ingenio 2010: MIPRCV (CSD2007-00018) and the Ministerio de Educacion y Ciencia under contract TIN2007-65533.

$\mathbf{Y}_i$ corresponds to the centroid of a Voronoi region.

2.1. Consistency with observation

Given an observed $\mathbf{Y}$, we can define the consistency set $R(\mathbf{Y}, \mathbf{P}, d)$ as all those images $\mathbf{X}$ where the colors at each pixel $\mathbf{X}_i$ belongs to the Voronoi region (defined by palette $\mathbf{P}$ and distance d) centered at $\mathbf{Y}_i$:

$$R(\mathbf{Y}, \mathbf{P}, d) = \{\mathbf{X} : \mathbf{X}_i \in V_{\mathbf{Y}_i}\ \forall\, i = 1, \ldots, N\}.$$

Our observation model is then based on giving uniform probability to images within the consistency set, and no probability otherwise. This leads to the following degenerate distribution:

$$p(\mathbf{Y}|\mathbf{X}) = \begin{cases} K & \text{if } \mathbf{Y} = f(\mathbf{X}) \\ 0 & \text{otherwise.} \end{cases}$$

where K is a real positive number.

2.2. Sparsity-based prior model

Let $\mathbf{\Phi}^t$ be a full column rank $M \times N$ matrix with $M > N$. Here, $\mathbf{\Phi}^t$ and $\mathbf{\Phi}$ represent, respectively, the analysis and synthesis operation of an overcomplete Parseval frame ($\mathbf{\Phi}\mathbf{\Phi}^t = \mathbf{I}$). In this work, we assume that each color channel is independent from each other. We note $\mathbf{x}_\zeta$, where $\zeta \in \{R, G, B\}$, to the corresponding $N \times 1$ RGB channel of image $\mathbf{X}$. Our image prior is based on expecting that each $\mathbf{\Phi}^t\mathbf{x}_\zeta = \mathbf{a}_\zeta + \mathbf{r}_\zeta$, where $\mathbf{a}_\zeta$ are sparse vectors and $\mathbf{r}_\zeta$ Gaussian correction terms. Then, by concatenating the column vectors $\mathbf{a}_\zeta$ into $M \times 3$ matrix $\mathbf{A}$, we define the following prior distributions:

$$\begin{aligned} p(\mathbf{X}|\mathbf{A}) &\propto \prod_{\zeta \in \{R,G,B\}} \exp\Big(\frac{-1}{2\sigma_{\mathbf{r}}^2}\|\mathbf{\Phi}^t\mathbf{x}_\zeta - \mathbf{a}_\zeta\|_2^2\Big), \\ p(\mathbf{A}) &\propto \prod_{\zeta \in \{R,G,B\}} \exp(-\gamma\|\mathbf{a}_\zeta\|_0). \end{aligned}$$

It is important to note that we can define in a similar way our signal model in any other orthogonal color space.

2.3. Problem formulation

We want to maximize the posterior probability $p(\mathbf{X}, \mathbf{A}|\mathbf{Y})$. By Bayes rule, this is equivalent to maximize the joint probability $p(\mathbf{X}, \mathbf{A}, \mathbf{Y})$, which we factorize as $p(\mathbf{X}, \mathbf{A}, \mathbf{Y}) = p(\mathbf{Y}|\mathbf{X})p(\mathbf{X}|\mathbf{A})p(\mathbf{A})$. Then, by minimizing the negative logarithm of $p(\mathbf{X}, \mathbf{A}, \mathbf{Y})$, we want to solve:

$$(\hat{\mathbf{X}}, \hat{\mathbf{A}}) = \arg\min_{\mathbf{X},\mathbf{A}} \Big\{ \sum_{\zeta \in \{R,G,B\}} \big(\|\mathbf{a}_\zeta\|_0 + \lambda\|\mathbf{\Phi}^t\mathbf{x}_\zeta - \mathbf{a}_\zeta\|_2^2\big) \quad \text{s.t. } \mathbf{Y} = f(\mathbf{X}) \Big\}, \tag{1}$$

where $\lambda = \frac{1}{2\sigma_{\mathbf{r}}^2\gamma}$.

3. A BLOCK COORDINATE DESCENT APPROACH

It is easy to see that minimizing the cost function of Eq. (1) in $\mathbf{A}$ for a given $\mathbf{X}$ results in the hard-thresholding of $\mathbf{\Phi}^t\mathbf{x}_\zeta$, for each ζ, with threshold $\theta = \lambda^{-\frac{1}{2}}$ (e.g. [9, 10]), which consists of setting to zero all coefficients in $\mathbf{\Phi}^t\mathbf{x}$ whose amplitude is below the threshold. We note $\mathbf{S}(\theta)\mathbf{\Phi}^t\mathbf{x}$ to this operation, where $\mathbf{S}(\theta)$ is a $M \times M$ diagonal matrix whose entries $S_{ii}(\theta) = 1$ if $\phi_i^t\mathbf{x} \geq \theta$ and $S_{ii}(\theta) = 0$ otherwise.

On the other hand, it can be proven that minimizing the cost function of Eq. (1) in $\mathbf{X}$ for a given $\mathbf{A}$ results in an orthogonal projection onto the consistency set.

But project onto Voronoi cells can be quite complicate, because we do not know, in general, onto which face of the polyhedron will fall the projection. We have approximated the projection by the following procedure. For each pixel outside its corresponding Voronoi cell, we do a linear search in the line between it and its corresponding centroid looking for the nearest point to the pixel which lies in the proper Voronoi cell, which can be checked by applying the quantization.

Finally, from an initial $\hat{\mathbf{X}}^{(0)}$ our method iterates as:

$$\begin{aligned} \mathbf{Z}^{(n)} &= \mathbf{\Phi}\mathbf{S}(\theta)\mathbf{\Phi}^t\hat{\mathbf{X}}^{(n)}, \\ \hat{\mathbf{X}}^{(n+1)} &= P^{\perp}_{R(\mathbf{Y})}(\mathbf{Z}^{(n)}). \end{aligned}$$

where $P_B^\perp$ is the orthogonal projection onto a set B. See section 4 for more details about the algorithm.

4. IMPLEMENTATION

Representation used. We have seen a good comparative performance by using the 8-scale Dual-Tree Complex Wavelet Transform (DT-CWT, [1]). The redundancy factor is 4. The real and imaginary part of complex coefficients are separated and treated as if they were independent coefficients. Both analysis and synthesis operations were performed on an orthogonal color domain obtained through matrix $T = [\frac{1}{\sqrt{3}}\ \frac{1}{\sqrt{3}}\ \frac{1}{\sqrt{3}};\ \frac{1}{\sqrt{2}}\ \frac{-1}{\sqrt{2}}\ 0;\ \frac{-1}{\sqrt{6}}\ \frac{-1}{\sqrt{6}}\ \frac{1}{\sqrt{6}}]$.

Algorithmic parameters. One of the critical issues of these kind of methods is choosing a value for parameter λ. During the last years, dynamic strategies based on beginning with a small value (equivalently, a high value of the threshold θ), which provides a high degree of sparsity, and increasing it at each iteration (decreasing θ) have provided excellent results in a wide variety of applications (e.g.,[7, 8, 10]). In [10], we justified the use of an exponential decay of θ for the sparse approximation case as a deterministic annealing-like strategy based on minimizing in successively less smoothed versions of the highly non-convex original cost function.

This is also the strategy used in this work. In the results below, we begin with and initial $\theta^{(0)}$ equal to the maximum amplitude of the coefficients $\mathbf{\Phi}^t\mathbf{y}_\zeta$, for $\zeta \in \{R, G, B\}$, and at each iteration n we multiply $\theta^{(n)}$ by $\beta = 0.8$. The iterations are stopped when current $\theta^{(n)}$ is below 1.

5. RESULTS AND DISCUSSION

Our test images are real gif images quantized with 256 colors. The images were downloaded from the web[1]. We compare here our proposed method vs. ℓ_1-norm minimization and constrained diffusion strategies. The former can be easily derived similarly to our ℓ_0-norm method, but using ℓ_1-norm in Eq. (1) instead. This leads to a method consisting of iterated dynamic soft-thresholding and projection onto consitency set. The algorithmic parameter setting is the same as for our proposed method. On the other hand, constrained diffusion represents the most popular approach to color de-quantizing in the literature. We have implemented a method consisting of convolving, at each iteration, the current estimate with $\mathbf{v}^t\mathbf{v}/256$, with $v = [1\ 4\ 6\ 4\ 1]$, and then projecting back to the consistency set. This is a much faster algorithm, and we saw that 15 iterations were enough to converge in all the cases.

Figure 1 shows an example where the quantization artifacts are strongly visible both in the smooth parts (sky) and in the textured areas (floor). Note that both ℓ_1-norm minimization and our proposed method are able to recover better the edges, which results in a more appropriate contrast of the trees with respect to the background and also of the textured areas. However, ℓ_1-norm minimization is not able to recover properly the original smoothness of the clouds, whereas constrained diffusion gives a too smooth result, loosing some details, like at the. Then, our method outperforms clearly the rest in this example.

Figure 2 shows another example, where strong artifact can be seen in the sky and in the crowd. Again, the result from constrained diffusion is too smooth and much texture is lost, whereas the result from ℓ_1-minimization is too sharp and the quantization artifact are not properly removed. Our method provides a more appropriate estimate, at the same time preserving well the edges due to the use of hard instead of soft thresholding but also removing the quantization artifacts even better than constrained diffusion.

6. CONCLUSIONS

We have proposed a new color de-quantizing method based on maximizing the compressibility of the response of overcomplete wavelets to images inside the consistency set with the degraded observation. Our observation model assumes uniform probability over the images within the consistency set, and zero probability outside. Our image model reflects the fact that most coefficients in the wavelet response are close to zero. Resulting method consists of iterated hard thresholding with a decreasing threshold at each iteration, followed by a projection onto the consistency set. We have provided some real examples showing that our method outperforms both ℓ_1-norm minimization and constrained diffusion tech-

[1] Unfortunately, the original URL seems to be now missing.

Fig. 1. From top to bottom: quantized observed image, result of constrained diffusion, result of ℓ_1-minimization, our result.

Fig. 2. From top to bottom: quantized observed image, result of constrained diffusion, result of ℓ_1-minimization, our result.

niques. More precisely, it is able to remove quantization artifacts effectively while preserving the sharpness of real edges.

7. REFERENCES

[1] N. Kingsbury, "Complex Wavelets for shift invariant Analysis and Filtering of Signals," *J. of App. and Comp. Harmonic Anal.*, **10** (3), pp. 234-253, May 2001.

[2] Z. Xiong, M.T. Orchard and Y. Zhang, "A deblocking algorithm for JPEG compressed images using overcomplete wavelet representations," *IEEE Trans. on Circuit and Syst. for Video Tech.*, **7** (2), pp. 443-437, April 1997.

[3] J. Mateos, A.K. Katsaggelos and R. Molina, "A Bayesian approach to estimate and transmit regularization parameters for reducing blocking artifacts," *IEEE Trans. on Im. Proc.*, **9** (7), pp. 1200-1215, July 2000.

[4] Y.H. Chan and Y.H. Fung, "A regularized constrained iterative restoration algorithm for restoring color-quantized images," *Signal Proc.*, **85** (7), pp. 13751387, July 2005.

[5] D. Keysers, C.H. Lambert and T.M. Breuel, "Color image dequantization by constrained diffusion, " in *Proceedings of the SPIE*, **6058**, January 2006.

[6] B.A. Olshausen, D.J. Field, "Natural Image Statistics and Efficient Coding," *Network Comp. in Neural Syst.*, **7**, pp. 333-339, 1996.

[7] J. Bobin, J.L. Starck, J. Fadili, Y. Moudden, and D.L. Donoho, "Morphological Component Analysis: an adaptive thresholding strategy," *IEEE Trans. on Im. Proc.*, **16** (11), pp. 2675–2681, November 2007.

[8] J. Portilla and L. Mancera, "L0-based sparse approximation: Two alternative methods and some applications, " in *SPIE Optics and Photonics*, San Diego, CA, August 26-30, 2007.

[9] P. Moulin and J. Liu, "Analysis of multiresolution image de-noising schemes using generalized-Gaussian and complexity priors," *IEEE Trans. on Inf. Theory, Special Issue on Multiscale Anal.*, **45** (3), pp. 909-919, April 1999.

[10] L. Mancera and J. Portilla, "Non-convex sparse optimization through deterministic anneling and applications, " in *IEEE Int. Conf. on Im. Proc.*, San Diego, CA, October 12-15, 2008.