Image De-Quantizing via Enforcing Sparseness in Overcomplete Representations

Luis Mancera and Javier Portilla*

Visual Information Processing Group,
Department of Computer Science and Artificial Inteligence,
Universidad de Granada
{mancera, javier}@decsai.ugr.es

Abstract. We describe a method for removing quantization artifacts (*de-quantizing*) in the image domain, by enforcing a high degree of sparseness in its representation with an overcomplete oriented pyramid. For this purpose we devise a linear operator that returns the minimum L2-norm image preserving a set of significant coefficients, and estimate the original by minimizing the cardinality of that subset, always ensuring that the result is compatible with the quantized observation. We implement this solution by alternated projections onto convex sets, and test it through simulations with a set of standard images. Results are highly satisfactory in terms of performance, robustness and efficiency.

1 Introduction

Spatial quantization is part of the image capture with digital devices. Usually artifacts (false contours and suppression of low-contrast texture) are close or even below the visibility threshold, but they become evident in a number of situations. For instance when stretching the local luminance range for detail inspection, or when de-convolving quantized blurred images, mostly if there is little random noise. It is also a useful step for local features extraction (e.g., luminance gradient) sensitive to those artifacts, for interpolating iso-level curves in topographic or barometric maps, or for using a reduced number of bits per pixel when there are not enough resources to perform image compression.

Surprisingly enough, de-quantizing in the image domain has received little attention in scientific literature (exceptions are [1,2]). In contrast, transform quantization has been widely treated, especially in the context of post-processing compressed images (de-blocking), usually under orthogonal or bi-orthogonal transforms (e.g., [3,4,5]), but also under overcomplete transforms (e.g., [6,7]). In this work we propose to enforce a certain characterization of sparseness in the overcomplete wavelet domain as the base criterion of the restoration, always ensuring that the estimated image is compatible with the quantized observation. The solution is formulated as belonging to the intersection of two convex sets [8].

* Both authors funded by grant TIC2003-1504 from the Ministerio de Ciencia y Tecnologia. JP is under the "Ramon y Cajal " program.

J. Blanc-Talon et al. (Eds.): ACIVS 2005, LNCS 3708, pp. 411–418, 2005.

2 Image Model

Linear representations based on multi-scale band-pass oriented filters (*wavelets*) are well-suited for representing basic properties of natural images, such as scale-invariance and the existence of locally oriented structures. Natural images typically produce *sparse* distributions of their wavelet coefficients. This means that the energy of the image is mostly concentrated in a small proportion of coefficients [9,10]. It has also been observed that overcomplete representations 1) may produce sparser distributions than critically sampled wavelets [11]; and 2) being translation invariant, they typically provide better results for image processing (e.g., [12,13]). For this work we chose the steerable pyramid [14], an oriented overcomplete representation whose basis functions are rotated and scaled versions of each other. In addition to marginal statistics in the wavelet domain, many authors have exploited the dependency existing among neighbor coefficients (e.g., [15,16,17,18]). We have considered this dependency when selecting significant coefficients, in Section 4.1.

3 Enforcing Sparseness

Most degradation sources decrease the sparseness of the wavelet coefficients (e.g., [19,20]). In an attempt to recover the high-sparseness condition of the original, we devise an operator which increases the image sparseness by preserving a given subset of *significant* coefficients while minimizing the global L2-norm. Now we describe this operator for overcomplete representations.

Let $\mathbf{x} \in \mathbb{R}^N$ be an image and $\mathbf{x}' = \mathbf{\Phi x}$ its overcomplete representation. $\mathbf{\Phi}$ is an $M \times N$ matrix ($M > N$) with each row ϕ_j representing an analysis function. We assume that $\mathbf{\Phi}$ preserves the L2-norm, $\|\mathbf{\Phi x}\| = \|\mathbf{x}\|$, and that it has perfect reconstruction, $\mathbf{\Psi \Phi x} = \mathbf{x}$, where $\mathbf{\Psi} = (\mathbf{\Phi}^T\mathbf{\Phi})^{-1}\mathbf{\Phi}^T$ is the pseudoinverse of $\mathbf{\Phi}$. Given an index set G of what we consider the M' most significant coefficients of $\mathbf{x}'$ (see Section 4.1), we define $\mathbf{\Phi}_G$ as the $M' \times N$ matrix formed by all ϕ_j row vectors such that $j \in G$. Our sparseness-enforcing operator for that set is:

$$\tilde{\mathbf{x}}(G, \mathbf{x}) = \arg \min_{\mathbf{z} \in \mathbb{R}^N} \|\mathbf{z}\| \text{ s.t. } \mathbf{\Phi}_G \mathbf{z} = \mathbf{\Phi}_G \mathbf{x}. \quad (1)$$

Naming $\mathbf{\Psi}_G$ the pseudoinverse of $\mathbf{\Phi}_G$, previous equation is equivalent to $\tilde{\mathbf{x}}(G, \mathbf{x}) = \mathbf{\Psi}_G \mathbf{\Phi}_G \mathbf{x}$. We call $\mathbf{S}_G = \mathbf{\Psi}_G \mathbf{\Phi}_G$. Note that, when $rank(\mathbf{\Phi}_G) = N$ the Equation (1) has a trivial solution, $\mathbf{x}$. Thus, for $\mathbf{S}_G$ to be of interest, we choose $M' < N$. However, note that $\mathbf{\Psi}_G$ is not trivial to compute. In practice we have applied the method of alternated projections (POCS), that states that iterative orthogonal projections onto a number of intersecting convex sets converge strongly to their intersection. We have used two convex sets: 1) the set of vectors of coefficients having the same values as $\mathbf{x}$ for the indices in G, $V(G, \mathbf{x}) = \{\mathbf{z}' \in \mathbb{R}^M : z'_j = \phi_j \mathbf{x}, \forall j \in G\}$; and 2) the set of admissible vectors of coefficients, $A(\mathbf{\Phi}) = \{\mathbf{z}' \in \mathbb{R}^M : \exists \mathbf{x}^0 \in \mathbb{R}^N : \mathbf{z}' = \mathbf{\Phi x}^0\}$. The orthogonal projection onto the first set is achieved by setting the coefficients with indices in G to their original

values, leaving the rest unchanged: $P^{\perp}_{V(G,\mathbf{x})}(\mathbf{z}') = \mathbf{D}(G)\mathbf{x}' + (\mathbf{I}(M) - \mathbf{D}(G))\mathbf{z}'$, where $\mathbf{x}' = \mathbf{\Phi}\mathbf{x}$, $\mathbf{I}(M)$ is the $M \times M$ identity matrix and $\mathbf{D}(G)$ is a $M \times M$ diagonal matrix such that $d_{ii} = 1$ if $i \in G$ and $d_{ii} = 0$ otherwise. For the other set, the orthogonal projection consists of inverting the transform and applying it again: $\mathbf{P}^{\perp}_{A(\mathbf{\Phi})}\mathbf{z}' = \mathbf{\Phi}\mathbf{\Psi}\mathbf{z}'$. The solution $\tilde{\mathbf{x}}(G, \mathbf{x})$ can be expressed as the inverse transform of the minimum L2-norm vector belonging to the intersection of both sets. So [1]: $\tilde{\mathbf{x}}(G, \mathbf{x}) = \mathbf{S}_G\mathbf{x} = \mathbf{\Psi} \lim_{n\to\infty}(\mathbf{P}^{\perp}_{A(\mathbf{\Phi})}P^{\perp}_{V(G,\mathbf{x})})^n(\mathbf{0})$. Since $P^{\perp}_{V(G,\mathbf{x})}(\cdot)$ is an affine orthogonal projector, superindex n indicates the number of iterative compositions of the functions within the brackets, not a power.

4 Image De-Quantizing

Left panel of Figure 1 shows a joint histogram of the coefficients for a subband of quantized *Boat* image vs. those of the original. It is normalized in amplitude by columns to express the probability of the observed given the original. We can see that low-amplitude coefficients are severely damaged whereas high-amplitude coefficients are just slightly damped and contaminated with noise of nearly constant variance. Right panel shows the same joint histogram, but normalized by rows and transposed, so now it expresses an empirical measurement of the posterior density of the original given the observation. We can also discriminate an inner region, for which the posterior density is complicated, and an outer region, where the original can be reliably estimated from the observation.

4.1 Selecting Significant Coefficients

According to the right panel of Figure 1, it seems natural to choose a threshold to discriminate the inner from the outer region. Discarding the effect of the prior density, it is also reasonable to choose, for each subband k, a threshold λ^{α}_k proportional to the standard deviation of the noise caused by the quantization in that subband, σ_k. That is, $\lambda^{\alpha}_k = \alpha\sigma_k$ with $\alpha \in \mathbb{R}^+$, where σ_k can be estimated analytically from the analysis functions, assuming that quantization noise is white and uniform in density. However, a more reliable estimate can be obtained through simulations, by measuring the variance for a set of quantized standard images, and averaging these measurements (our choice). Once we have estimated every σ_k for a certain quantization process, we can sweep the proportionality factor α from 0 to ∞ to control the cardinality of G, as a sparseness index of $\tilde{\mathbf{x}}(G, \mathbf{x})$. We have obtained better results for this task by considering that a coefficient is significant if *any* coefficient within its neighborhood, including itself, surpasses the amplitude threshold λ^{α}_k. For this work we have used a 5×5 spatial neighborhood. Naming I^0_k the set of indices of subband k and $N(\cdot)$ the neighborhood set for a given index, the set of significant coefficients of k is:

$$I_k(\alpha) = \{i \in I^0_k : i \in N(j); \forall j \in I^0_k : |\phi_j\mathbf{x}| \geq \alpha\sigma_k\},$$

and the total set of significant indices is: $G(\alpha) = \bigcup_{k=1}^{K} I_k(\alpha)$.

[1] We have greatly accelerated convergence by using a linear prediction technique.

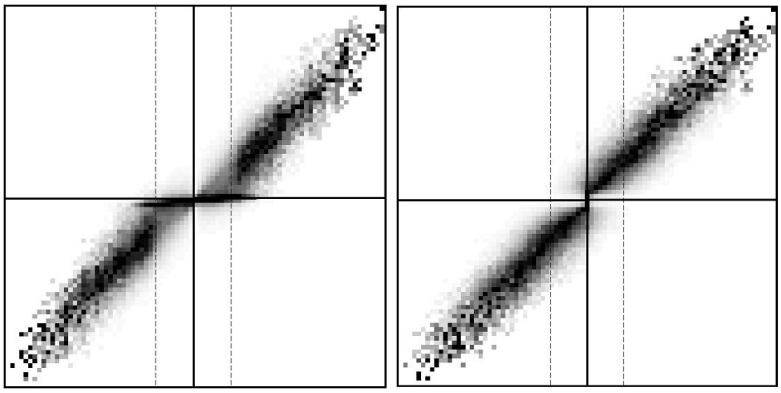

Fig. 1. Normalized joint histograms of the coefficients of a subband from a quantized image and from the original image. *Left*: degradation model (original in abscissas). *Right*: Posterior density (quantized in abscissas).

4.2 Global Problem Formulation

Let's consider an index set $G(\alpha)$ and let $q(\mathbf{x}) = \mathbf{y}$ be a quantization process. When applying our sparseness-enforcing operator, it exists the possibility that $q(\mathbf{S}_{G(\alpha)}\mathbf{x}) \neq \mathbf{y}$. We must ensure that the final estimation belongs to $Q(\mathbf{y}) = \{\mathbf{x} \in \mathbb{R}^N : q(\mathbf{x}) = \mathbf{y}\}$, the compatibility set for $\mathbf{y}$. On the other hand, it is easy to check that $\mathbf{S}_{G(\alpha)}$ is the orthogonal projector onto the set $C(G(\alpha)) = \{\mathbf{x} \in \mathbb{R}^N : \exists \mathbf{x}^0 \in \mathbb{R}^N \text{ s.t. } \mathbf{x} = \mathbf{S}_{G(\alpha)}\mathbf{x}^0\}$, which represents the set of *sparsified* images obtained with the set of indices $G(\alpha)$. This is a linear subspace. Its dimensionality is the cardinality of $G(\alpha)$, which decreases as α increases. As we are looking for the smallest possible set of significant coefficients, we search for the highest α such that $C(G(\alpha))$ still includes at least one image compatible with $\mathbf{y}$:

$$\hat{\alpha} = \sup\{\alpha \in \mathbb{R}^+ : C(G(\alpha)) \bigcap Q(\mathbf{y}) \neq \emptyset\}. \tag{2}$$

Calling $T(\alpha, \mathbf{y}) = C(G(\alpha)) \bigcap Q(\mathbf{y})$, we choose our final estimate to be the element of $T(\hat{\alpha}, \mathbf{y})$ closest to the observation $\mathbf{y}$: $\hat{\mathbf{x}} = P^{\perp}_{T(\hat{\alpha},\mathbf{y})}(\mathbf{y})$, where $P^{\perp}_{T(\hat{\alpha},\mathbf{y})}$ is a (non-linear) orthogonal projection function. In contrast with most estimators used in restoration, this preserves all the information carried by the observation.

4.3 POCS-Based Solution

As $Q(\mathbf{y})$ and $C(G(\alpha))$ are both convex sets, we can compute our estimation for a given α, noted $\hat{\mathbf{x}}^{\alpha}$, through alternated orthogonal projections. The orthogonal projection onto $Q(\mathbf{y})$ can be defined as $\mathbf{z} = P^{\perp}_{Q(\mathbf{y})}(\mathbf{x})$, with

$$z_i = \begin{cases} x_i, & y_i - \frac{\delta_i}{2} < x_i \leq y_i + \frac{\delta_i}{2} \\ y_i - \frac{\delta_i}{2} + \epsilon, & x_i \leq y_i - \frac{\delta_i}{2} \\ y_i + \frac{\delta_i}{2}, & y_i + \frac{\delta_i}{2} < x_i \end{cases}$$

where δ_i are each quantization interval width and $\epsilon \in \mathbb{R}^+$ (ideally infinitesimal) is an artifice to achieve empty intersection between adjacent closed intervals. Therefore, our estimation for a given α is: $\hat{\mathbf{x}}^\alpha = \lim_{n\to\infty}(P^{\perp}_{Q(\mathbf{y})}\mathbf{S}_{G(\alpha)})^n\mathbf{y}$. We use a line search to find the highest factor $\hat{\alpha}$ for which previous limit converges.

4.4 An Efficient Approximated Solution

We have verified two very positive facts. First, that the factor $\hat{\alpha}$ closely match the LSE hand-optimized factor in simulations. Second, that $\hat{\alpha}$ is remarkably constant for different images (typically ranging between 4 and 5), under the same linear representation and the same quantization. Thus we have used, instead of a different $\hat{\alpha}$ each time, an averaged value computed off-line for a collection of standard images. Doing that, we save around one order of magnitude in computation time (which becomes close to 10 s. for 256^2 images and to 50 s. for 512^2, with our MATLAB implementation using a 3.4 Ghz Pentium IV CPU), whereas decrease in Signal-to-Noise Ratio is only around 0.10 dB. We note that when the average factor is higher than the optimal value, $Q(\mathbf{y})$ and $C(G(\alpha))$ do not intersect, and then POCS provides a LS-optimal solution.

5 Results and Discussion

We have tested our method on a set of ten 8-bit grayscale standard images, of 256^2 pixels (*Lena, Peppers, Cameraman, Einstein, House*) and 512^2 (*Barbara, Boat, Goldhill, Plane, Windmill*). We have used the steerable pyramid [14] with 6 scales and 4 orientations. Figure 2 shows the increase in Signal-to-Noise Ratio (ISNR) of the processed images w.r.t. the observations for a range of quantization bits. Improvement is remarkable, especially in the medium range. There is a sudden descent in performance in the fine quantization range that we have not completely explained yet.

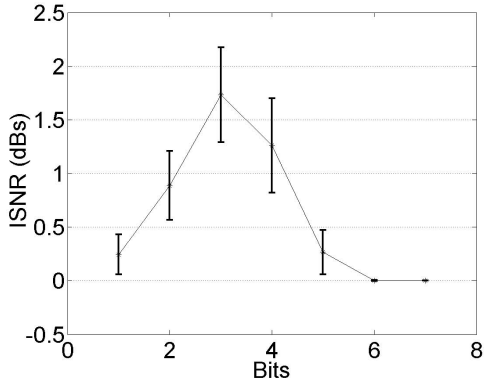

Fig. 2. Increment in SNR (ISNR), expressed as $10 \cdot log_{10}(\sigma_q^2/\sigma_r^2)$, where σ_q^2 and σ_r^2 are the MSE for observation and estimation, respect., and for several quantization bits

Table 1. Results of our method for 3 and 4 bit quantization, showing the increment w.r.t. the original in Signal-to-Noise-Ratio (in dB) and also in the Structural Similarity Index (×100) [21]

Bits	Metric	Ref.	*Barbara*	*Boat*	*Lena*	*Peppers*
3	ISNR	28.74	2.29	2.24	1.93	2.01
	ISSIM	80.53	7.29	6.03	6.71	7.48
4	ISNR	34.77	1.91	0.86	1.54	1.38
	ISSIM	90.10	3.70	0.85	3.24	3.05

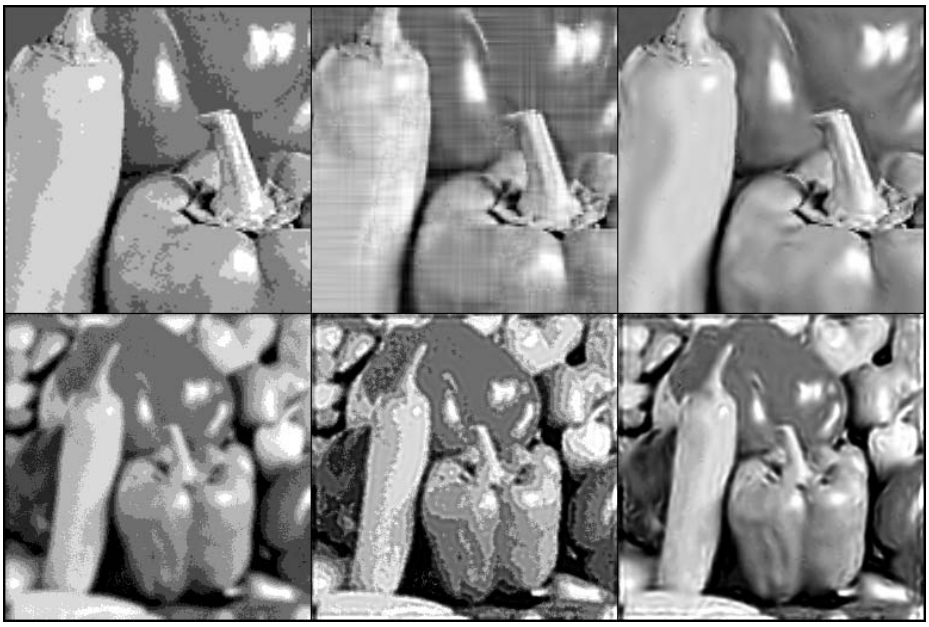

Fig. 3. Some visual results. *Left*: Observation. *Center*: Comparison method. *Right*: Our method. See text for details.

Table 1 shows numerical results for quantization using 3 and 4-bits. We also include the increase in the Structural Similarity Index (SSIM) [21], a perceptually-inspired metric taking values in the range [0,1]. First column shows the averaged PSNR/SSIM values of the observation (SSIM values multiplied times 100). There is a very significant improvement under both metrics. First row of Figure 3 shows a cropped result using 3 bits. We have implemented a method (central panel) to help us as a reference, based on a similar strategy than [2]: a gradient-descent in the L2-norm of the output of a high-pass (Laplacian) filter applied to the image, each time projecting the updated image onto the compatibility set $Q(\mathbf{y})$. Our result (right panel) is 1.30 dB above, confirming once again that sparseness-based solutions are more powerful than classical

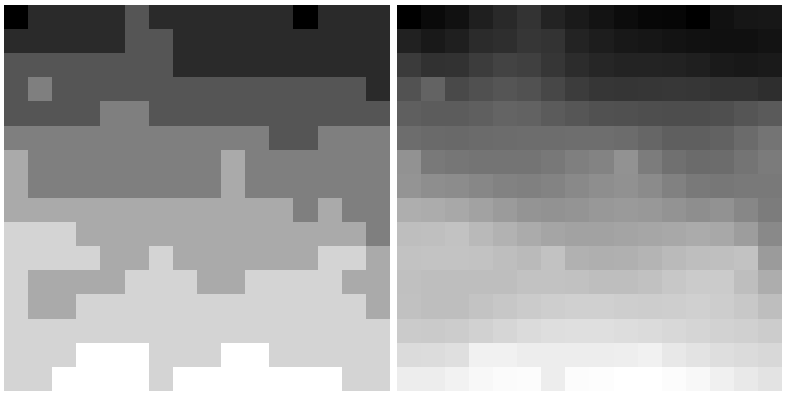

Fig. 4. *Left.* Detail of the sky of a real 8-bits photographic picture (amplified contrast). *Right.* Same detail in the picture after processed by our method.

smoothness-based approaches. Second row shows another example. The same image is blurred with a Gaussian kernel ($\sigma_b = \sqrt{2}$), corrupted with noise ($\sigma_n = 2$) and then quantified with 3 bits. Central panel shows the results from a general purpose maximum-likelihood semi-blind deconvolution method (**deconvblind** in MATLAB), passing the Point Spread Function as argument but not the variance of the noise. Right panel shows the deconvolution after de-quantizing the observation. The suppression of artifacts is very noticeable (1.18 dBs ISNR, 0.13 ISSIM). It is remarkable that we have obtained for the image in the right panel the same SSIM (0.70) w.r.t. the original as for the result of applying the deconvolution directly to the unquantized blurred and noisy image. Figure 4 shows removal of low-contrast artifacts in a 8-bit image.

To conclude, we have presented an automatic sparseness-based practical tool for removing pixel quantization artifacts which provides close to LS-optimal results. We still have to further investigate the causes for the drop in performance on the fine quantization range.

References

1. Desolneux, A., Ladjal, S., Moisan, L., Morel, J.M.: Dequantizing Image Orientation. IEEE Trans. Image Proc. **11**, 10, (Oct 2002) pp. 1129-1140.
2. Chan, Y.H., Fung, Y.H.: A Regularized Constrained Iterative Restoration Algorithm for Restoring Color-Quantized Images. Elsevier Sig. Proc., **85**, pp. 1375-1387, 2005.
3. Paek, H., Kim, R., Lee, S.: On the POCS-based Postprocessing Technique to Reduce the Blocking Artifacts in Transform Coded Images. IEEE Trans. Circuit and Syst. for Video Tech., **8**, 3, (Jun 1998), pp. 358-367.
4. Mateos, J., Katsaggelos, A.K., Molina, R.: A Bayesian Approach to Estimate and Transmit Regularization Parameters for Reducing Blocking Artifacts, IEEE Trans. Image Proc., **9**, 7, (2000), pp. 1200-1215.

5. Li, X.: Improved Wavelet Decoding via Set Theoretic Estimation. Images. IEEE Trans. Circuit and Syst. for Video Tech., **15**, 1, (Jan 2005), pp. 108-112.
6. Xiong, Z., Orchard, M.T., Zhang, Y.: A Deblocking Algorithm for JPEG Compressed Images Using Overcomplete Wavelet Representations. IEEE Trans. Circuit Syst. Video Tech., **7**, 2, (Apr 1997), pp. 433-437.
7. Goyal, V.K., Vetterli, M., Thao, N.T.: Quantized Overcomplete Expansions in $\mathbb{R}^n$: Analysis, Synthesis and Algorithms. IEEE Trans. Inf. Theory, **44**, 1, (1998), pp.16-31.
8. Youla, D.C.: Generalized Image Restoration by the Method of Alternating Orthogonal Projections. IEEE Trans. Circuits and Syst., **CAS-25**, 9, (Sep 1978).
9. Mallat, S.G.: A Theory for multiresolution signal decomposition: The wavelet representation. PAMI, **11**, (Jul 1989, pp. 674-693.
10. Olshausen, B.A., Field, D.J.: Natural Image Statistics and Efficient Coding. Network Computation in Neural Systems, **7**, (1996), pp. 333-339.
11. Olshausen, B.A., Field, D.J.: Sparse Coding with an Overcomplete Basis Set: A Strategy Employed by V1?. Vision Res., **37**, 23, (1997), pp. 3311-3325.
12. Simoncelli, E.P., Freeman, W.T., Adelson, E.H., Heeger, D.J.: Shiftable Multi-Scale Transforms. IEEE Trans. Inf. Theory, **38**, 2, (Mar 1992), pp. 587-607.
13. Coifman, R.R., Donoho, D.L.: Translation Invariant De-noising. Lecture Notes in Statistics, **103** (1995) pp. 125-150.
14. Simoncelli, E.P.: The Steerable Pyramid: A Flexible Architecture For Multi-Scale Derivative Computation. 2nd IEEE Intl Conf. Im. Proc., **III**, pp.444-447, Oct 1995.
15. Shapiro, J.: Embedded Image Coding Using Zerotrees of Wavelet Coefficients. IEEE Trans. Signal Proc., **41**, 22, (Dec 1993), pp. 3445-3462.
16. Buccigrossi, R.W., Simoncelli, E.P.: Image Compression via Joint Statistical Characterization in the Wavelet Domain. IEEE Trans. Image Proc. **8**, 12, Dec 1999, pp. 1668-1701.
17. Pižurica, A., Philips, W., Lemahieu, I., Acheroy, M.: A joint inter- and intrascale statistical model for Bayesian wavelet based image denoising. IEEE Trans. Image Proc., **11**, 5, (May 2002), pp. 545-557.
18. Portilla, J., Strela, V., Wainwright, M.J., Simoncelli, E.P.: Image Denoising using Scale Mixtures of Gaussians in the Wavelet Domain. IEEE Trans. Image Proc., **12**, 11, (Nov 2003), pp. 1338-1351.
19. Rooms, F., Philips, W., Portilla, J.: Parametric PSF estimation via sparseness maximization in the wavelet domain. SPIE Conference "Wavelet Applications in Industrial Processing II" (Oct 2004), Philadelphia. Proc. SPIE **5607**, pp. 26–33.
20. Wang, Z., Wu, G., Sheikh, H.R., Simoncelli, E.P., Yang, E.H., Bovik, A.C.: Quality-Aware Images. IEEE Trans. on Image Proc., accepted, 2005.
21. Wang, Z., Bovik, A.C., Simoncelli, E.P.: Image Quality Assessment: from Error Visibility to Structural Similarity. IEEE Trans. Im. Proc. **13**, April 2004,pp.600-612.