

Comparing Objective and Subjective Quality Results for Compression Pre-processing with Non-linear Diffusion

Ivan Kopilovic¹, Tamás Szirányi²

¹ University of Konstanz, Department of Computer & Information Science, Fach M 697
D-78457 Konstanz, Germany

kopilovi@inf.uni-konstanz.de

² Analogical and Neural Computing Laboratory, Comp. & Automation Inst., Hungarian Academy of Sciences,

H-1111 Budapest, Kende u. 13-17., Hungary

and

University of Veszprém, Department of Image Processing and Neurocomputing
sziranyi@sztaki.hu

Abstract. Compression systems like JPEG include optional pre-processing with filtering to avoid compression artefacts. At higher compression ratios a stronger filtering is needed that impacts the large scale image content. To preserve the large scale information we have previously proposed to use non-linear diffusion as a pre-processing for filtering out small scale details irrelevant at a given compression ratio and acting as noise. Now we compare typical diffusion processes applied before the blockwise DCT compression using the peak signal to noise ratio (PSNR) as an objective quality measure. We give a simple measure of artefact reduction in terms of PSNR, and show that a considerable artefact reduction is achieved by pre-processing at the same bit rate as and with no greater error than the original compression. We did tests to see if the above artefact reduction implies a better subjective impression of quality. The images processed with the PSNR-based algorithm had nearly the same but greater PSNR value as the original compression. Subjects preferred noisy image content to the lack of small scale details, so the subjective preference of the images with reduced artefact is worse than of the original compression. Results suggest however that non-linear diffusion is more efficient for artefact reduction than non-adaptive smoothing like Gaussian filtering in terms of the subjective preference.

1 Introduction

Lossy image and video compression yield typical error patterns on the decompressed images or video sequences due to the quantisation error. Depending on the compression scheme and the bit rate, these can be ringing patterns around the edges, false or blurred texture, visible block-boundaries in block-partitioning schemes [20]. These phenomena are called compression artefacts. Compression artefacts not only deterio-

rate the visual quality of images but also perturb image processing and machine vision algorithms like edge detection and motion estimation.

A possible solution is to do a post-processing on the decompressed image, which removes compression artefacts. Some post-processing methods improve the visual quality of the compressed image by various adaptive filters that alleviate the artefacts and enhance the edges [3,4,9,19,20]. The blocking artefact for the blockwise transform compression can also be reduced by prediction methods for the quantised or missing transform coefficients [14]. Post-processing can also be formulated as a reconstruction problem for the original image and solved by optimisation with respect to an objective error measure like MSE, to probability considerations, to regularity constraints, or to a combination of these [11,17, 24].

The quality of the compressed images is affected not only by the bit rate and the compression algorithm, but also by the data source. In particular, images with more details usually degrade more than those with fewer details when compressed at the same bit rate [20]. It is therefore sensible to do some pre-processing on the image before compression, which alleviates the components of the image susceptible to artefacts. Pre-filtering methods were proposed for video coding [15]. Although pre-filtering options for still images exist in compression software products (e.g. in [25]), little research has been done in this area. Pre-processing has several advantages. It does not require a change in the decompression, and the compression standard remains unaffected. Moreover, pre-processing is done only once, while any post-processing must be done each time the image is decompressed, involving additional computational complexity. As a practical application, digital cameras use motion-JPEG for storing images and involve some kind of pre-processing compensating for the different optical and/or compression artefacts.

In this paper, we consider the idea in [21] in a systematic way. We define what we mean by artefact reduction, and use a simple way of measuring and expressing it through PSNR. Relying on the theory of anisotropic diffusion filtering, we claim that it is possible to achieve artefact reduction, while preserving the main structure of the image. To check whether artefact reduction results in images that are more preferred by human observers, subjective tests were done to rate the different methods. Results on test images suggest that the block boundaries and false textures are alleviated in the pre-processed images. The diffusion strength is controlled by PSNR measurements. The results of the subjective tests show however that the maximal artefact reduction obtained this way does not increase the subjective preference of the compressed image. This follows from the over-smoothing of some textured regions. The subjects prefer compressed images with noisy texture to images with reduced noise but less small-scale details. The tests show that there are significant differences between the diffusion methods. In particular, the more edge-adaptive the method is in the mathematical sense, the more it will be preferred by the observers.

2 Non-linear Diffusion and Adaptive Filtering

We consider diffusions that are potentially useful for artefact reduction having different levels of edge-adaptability. We begin by analysing the *linear diffusion* (LD) process.

We treat images as positively valued smooth functions defined at the points $\mathbf{x} \in \mathfrak{R}^2$. Let us take the smoothing of an image f with Gaussian kernels

$$G_t(\mathbf{x}) = \frac{1}{4\pi t} \exp\left(-\frac{|\mathbf{x}|^2}{4t}\right),$$

$\mathbf{x} \in \mathfrak{R}^2$ having various parameters $t > 0$. We obtain a family of images $(u_t)_{t \geq 0}$ with $u_0 = f$ and $u_t = G_t * f$. Each element u_t of this family can also be obtained [6,10] by an LD process done on the image up to time t , i.e., $u_t(\mathbf{x}) = u(t, \mathbf{x})$ for all $t > 0$ and $\mathbf{x} \in \mathfrak{R}^2$, where u is the solution of the LD equation

$$\partial_t u = \Delta u, \quad (1)$$

with the initial condition $u_0 = f$. The LD generates a multiscale representations of images [1,6], where t is called the scale of the diffusion.

To understand the usefulness of this filtering for artefact reduction, the Laplacian operator Δ is written as a sum of two orthogonal components

$$\Delta u = u_{\parallel} + u_{\perp}, \quad (2)$$

where u_{\parallel} denotes the second spatial derivative in the direction orthogonal to the gradient ∇u , and u_{\perp} is the second spatial derivative in the direction parallel with the gradient ∇u . The term u_{\parallel} can be interpreted as an ‘‘infinitesimal’’ Gaussian filtering along the edge, and u_{\perp} as an ‘‘infinitesimal’’ Gaussian filtering across the edge. This low-pass filtering can contribute to the reduction of the ringing artefact, which is the result of the sharp frequency cut-off caused by quantisation at lower bit rates. In Eq. 2, the edge-parallel and edge-normal directions have equal weights, and both diffusion terms depend only on the direction of ∇u and not on $|\nabla u|$, where the latter indicates the local contrast difference. The smoothing obtained by the LD respects neither the direction nor the contrast of the edges.

To add contrast and directional sensitivity we extend Eq. 2 as

$$\partial_t u = p\left(|\nabla G_{\sigma} * u|\right) u_{\parallel} + n\left(|\nabla G_{\sigma} * u|\right) u_{\perp}, \quad (3)$$

where p and n are weighting functions controlling the diffusion along and across the edges, respectively, and $\sigma > 0$ is a fixed parameter. The purpose of the pre-smoothing with G_{σ} is to obtain a reliable estimate on edges, and to make the equation robust against noise. We allow full diffusion at uniform regions where the value of $|\nabla G_{\sigma} * u|$

Table 1 Diffusions with different degree of adaptability.

	<i>Directionally insensitive (isotropic)</i> $\alpha = 0.5$	<i>Directionally adaptive (anisotropic)</i> $\alpha = 0$
<i>Contrast insensitive</i> $K = \infty$	Linear Diffusion (LD)	Mean Curvature Motion Diffusion (MCMD) [1]
<i>Contrast adaptive</i> $0 < K < \infty$	Non-linear Isotropic Diffusion (NLID) [18]	Pure Anisotropic Diffusion (PAD) [2,5]

is small and to inhibit the diffusion at edge locations where $|\nabla G_\sigma * u|$ is large. One possibility to control the diffusion in this way is to use the weighting function

$$w_K(x) = \begin{cases} 2 \exp\left(-\frac{|x|^2}{K}\right) & , 0 < K < \infty \\ 2 & , K = \infty \end{cases}, \quad (4)$$

where $x \in \mathfrak{R}$ and $K \in (0, +\infty]$ is a fixed parameter [16]. With the special choice $p = (1 - \alpha)w$ and $n = \alpha w$, where $0 \leq \alpha \leq 0.5$, we obtain the diffusions examined in this paper

$$\partial_t u = w_K(|\nabla G_\sigma * u|) \left((1 - \alpha)u_{\parallel} + \alpha u_{\perp} \right), \quad (5)$$

The parameter α controls the directional adaptability and the parameter K controls the contrast sensitivity. The diffusions obtained for particular choices of the parameters are listed in Table 1.

Whatever the parameters α and K are, the diffusion can contribute to the suppression of the ringing artefact to some degree, as explained above. Moreover, since the minima and the maxima of the intensity values get closer (the contrast decreases), the DC values of the neighbouring DCT blocks of a flat area are more likely to fall into the same quantisation bin after the pre-processing, thus decreasing the blocking artefact.

We used forward numerical schemes for the above equations in all of our experiments with the fixed step size $\lambda = 0.1$, and parameters $K = 0.05$, $\sigma = 0.4$, by recasting $(1 - \alpha)u_{\parallel} + \alpha u_{\perp} = \alpha \Delta u + (1 - 2\alpha)u_{\parallel}$ and computing u_{\parallel} as in [2]. The smoothing with G_σ for the gradient computation was done with LD. The scale values whenever indicated were computed as $t = \lambda m$, where m is the number of iterations done by the numerical scheme.

3. Artefact Reduction

To explain artefact reduction, we observe the details of the different compressed versions of the test image “Goldhill” shown in Fig. 3. The JPEG compressed image clearly suffers from ringing and blocking artefacts. By diffusion pre-filtering the artefacts will be reduced. It can be seen in Figs. 3b-d, if we compare the images along the high-contrast edges, or surfaces where block-boundaries become less visible.

Apart from visual quality, artefacts can also impair the stability of a computer vision algorithm. As an example, we give the edge detection results for the compressed versions of the image “Boat” from Fig. 4. The contours were extracted with a Canny edge detector [12] using fixed parameters. We compared the edges with the edges of the original image, and separated the true edge points from the falsely detected edge points. The true edge points are shown in Figs. 5 and the false detections are shown Fig. 6. Clearly, though the compressed image without pre-processing lets more original edge points detected (mainly texture), it will also give rise to more false structures. The ration of the number of true and false edge points in this example is 1.38 for the compressed image and 1.68 for the compressed image with pre-processing.

We try to define and measure artefact reduction first. Let P_t denote a diffusion pre-processing method done up to scale t and C the compression method (JPEG in our case). We assume for now that the bit rate is fixed. We say that an image f can be compressed at better quality than an image g at the given rate if $PSNR(C(f), f) \geq PSNR(C(g), g)$, where $C(f)$ and $C(g)$ denote the compressed images, and $PSNR(f, g) = -10 \log(255^2 / MSE(f, g))$, with N denoting the image size, and MSE the mean square error. This quality measure is used usually to evaluate compression results. Our goal is to transform f into an image f' by pre-processing so that we can compress f' at a better quality than f .

Suppose that we pre-process the image f with a diffusion method P up to scale t . The reconstruction quality of the pre-processed image Pf at the given fixed bit rate will be

$$Q_{PP}(t) = PSNR(C(P_t f), P_t f).$$

The quality of the compression with pre-processing relative to the original image f is

$$Q_P(t) = PSNR(C(P_t f), f).$$

The reconstruction quality of the compression without pre-processing is

$$Q_0 = Q_{PP}(0) = Q_P(0) = PSNR(C(f), f).$$

A diffusion processing method is said to reduce artefact if there is a scale $t \geq 0$ such that Pf can be compressed at a better quality than f , i.e., if $Q_{PP}(t) > Q_0$, under the constraint $Q_P(t) \geq Q_0$. We impose the latter constraint, since any processing is meaningful only if the original quality does not decrease.

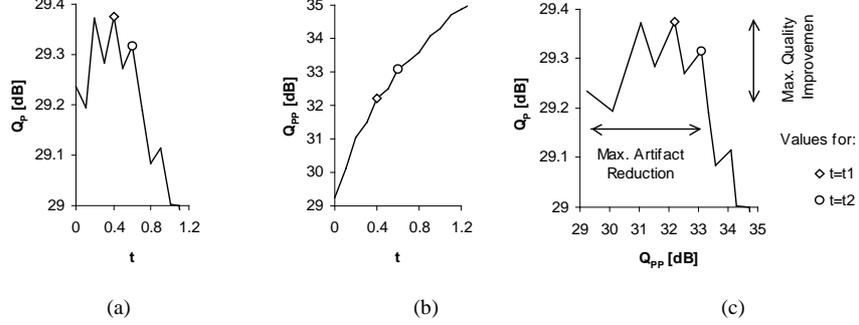


Figure 1 Pre-processing for “Goldhill” with NLID ($c = 0.25$ bits/pixel). (a) the PSNR values v. the original image $Q_P(t) = PSNR(C(P_t f), f)$, (b) the PSNR values v. the processed image $Q_{PP}(t) = PSNR(C(P_t f), P_t f)$, (c) the curve $t \mapsto (Q_{PP}(t), Q_P(t))$.

The tendencies of the curves Q_P and Q_{PP} as a function of t are shown for “Goldhill” in Fig. 1a and 1b for the pre-processing with NLID at bit rate $c = 0.25$ bits/pixel. The combined plot of $Q_P(t)$ and $Q_{PP}(t)$ as a function of the increasing scale $t \geq 0$ is shown in Fig. 1c. Important to note is that the value $Q_{PP}(t)$ increases monotonously with t , and that $Q_P(t)$ changes around the value Q_0 , reaches a maximum, and drops for larger scales. We can use $Q_{PP}(t)$ to measure the artefact reduction.

The latter observations and our definition of the artefact reduction lead us to the definition of the following two characteristic scales for each particular pre-processing method P :

1. The scale corresponding to the largest artefact reduction with maximal quality improvement,

$$t_1 = \max \arg \max \{ Q_P(t) \mid t \geq 0 \}. \quad (6)$$

2. The scale corresponding to the maximal artefact reduction,

$$t_2 = \max \{ t \geq 0 \mid Q_P(t) \geq Q_0 \}. \quad (7)$$

These scale values are indicated in Fig. 1, t_1 with a diamond and t_2 with a circle. Note that the pre-processing up to scale t_2 will redistribute the original compression error, so that the maximum portion of the quality will be devoted to the main structure of the image, as defined by the underlying multiscale representation, and the smallest portion of quality will be allotted to the small-scale details and noise. Different diffusion methods will do this redistribution in different ways. According to the theory concerning the filtering with non-linear diffusion, the more adaptive diffusion is, the better this redistribution will be. The results of the subjective tests support this claim.

Table 2 The PSNR values Q_P (compression quality) and Q_{PP} (artefact reduction) for the maximal quality improvement (scale t_1) and the maximal artefact reduction (scale t_2). Q_0 is the original JPEG compression quality in PSNR (dB).

		0.25 bits/pixel			0.4 bits/pixel		
		Goldhill	Boat	Bridge	Goldhill	Boat	Bridge
Q_0		29.23	30.11	24.09	30.87	32.40	25.39
$Q_P(t_1)$	LD	29.36	30.26	24.37	30.91	32.43	25.67
	NLID	29.37	30.32	24.41	30.97	32.46	25.61
	PAD	29.41	30.37	24.42	31.02	32.56	25.67
$Q_{PP}(t_1)$	LD	31.15	31.77	26.32	33.18	34.53	28.12
	NLID	32.19	32.28	25.74	34.00	33.28	27.90
	PAD	30.46	31.00	25.31	32.33	33.43	27.20
$Q_P(t_2)$	LD	29.32	30.12	24.24	30.91	32.43	25.48
	NLID	29.32	30.13	24.12	30.97	32.43	25.44
	PAD	29.24	30.14	24.09	30.87	32.40	25.43
$Q_{PP}(t_2)$	LD	32.52	32.99	29.98	33.18	34.53	29.86
	NLID	33.10	33.13	29.21	34.00	33.98	29.74
	PAD	32.70	32.15	28.70	33.23	34.01	28.51

4. Results

We compare the pre-processing with three diffusion processes (LD, NLID, PAD) for three typical test images at two selected bit rates. The results are summarized in Table 2. The table contains the original compression quality Q_0 without pre-processing, the maximal improvement in quality $Q_P(t_1)$ with the corresponding value $Q_{PP}(t_1)$ indicating the artefact reduction, and the quality values $Q_P(t_2)$ and $Q_{PP}(t_2)$ corresponding to the maximal artefact reduction (compare Fig. 1). Note that the value $Q_P(t_2)$ is very close, but in general it is not equal to Q_0 . This is because we cannot change the scale parameter and the quantisation intervals continuously.

These values show that it is possible to improve the compression quality in PSNR up to 1.5%, depending on the bit rate. As regards the artefact reduction, the pre-processed images can be compressed at 5 - 10% better quality than the original image at the pre-processing scale t_1 , and at up to 20% better quality at the pre-processing scale t_2 , without loss in the quality as compared to the original image (Q_0). The relative improvement in any sense is better for the lower bit rate.

If we look at the images in Fig. 3, we can see that the pre-processed images have less artefact. The LD-pre-processed image has got more blurred than those pre-processed with non-linear diffusion.

Though we have shown that it is possible to obtain a maximal improvement in the PSNR compression quality by pre-processing up to scale t_1 , this will not yield a large perceivable artefact reduction in general. With the maximal artefact reduction, there is a larger perceivable artefact reduction, but in spite of the PSNR results, which suggest improvement, we should test whether the subjective quality is deteriorated by the blur involved in these diffusion processes.

For the above mentioned reasons, we were interested in comparing the different pre-processing methods on a subjective scale for the maximal artefact reduction (scale t_2). We did a so-called Thurstone-scaling experiment [22,23]. The subjective scale is constructed based on a larger number of pairwise comparisons of the different stimuli. Thurstone-scaling is typically used in situations where the stimuli are very similar (in our case, we have nearly equal MSE for the images) or where the measured quantity is hard to describe exactly (image quality is hard to define in general).

In the scaling experiment, a test image is taken, and the test person has to compare each version of the image with each other version of the image. The versions constitute the compressed image without pre-processing, and the compressed images with the different pre-processing methods (only one pair is shown on the display at a time). We made the test for three different images. The scaling method gives useful results if the stimuli are very similar and do not contain outliers. For these reasons, the original images were excluded from the comparisons. *In this way, we did not measure the similarity (distance) to the original image, but a subjective impression of quality manifesting itself in the preference decisions of the human observers.*

We show the results in Fig. 2 obtained for a subjective test with 21 test persons. We constructed scales for the two compression rates by collecting the scores for the different methods across the different test images. The scales for these overall scores are also shown in Fig. 2. The scale is interpreted roughly as follows: the guesses of subjects for each item (JPEG, LD, etc.) are modelled as normally distributed random variables having the same variance, which is equal to the unit of the scale. The scale value for each item is the expected value of the corresponding random variable. The unit of the scale is interpreted as the uncertainty of the guesses of subjects. We emphasize that the scales in Fig. 2 are not directly comparable, since the units are obtained implicitly through the construction of each particular scale, and depend on the image.

The parameters of the above Gaussian distributions are computed based on the preference probabilities for each pair of images (the probability that one image is preferred to the other). These probabilities are in turn obtained by counting the scores for each pair. There are more possibilities for measuring the reliability of the scales. The χ^2 -test of Mosteller [13] is typically used for these purposes. Another way to test reliability is to compute the regression coefficient between the set of preference probabilities and the preference probabilities recomputed with the parameters of the Gaussian model. The results of the reliability test are summarized in Table 3. Because the scale for "Boat" at 0.4 *bits/pixels* is not very reliable, the scores for this case were not considered in the construction of the overall scale for 0.4 *bits/pixels*.

Although the compared images have almost the same PSNR error (quality values $Q_P(t_2)$ in Table 2), their preference by the subjects is quite different. The

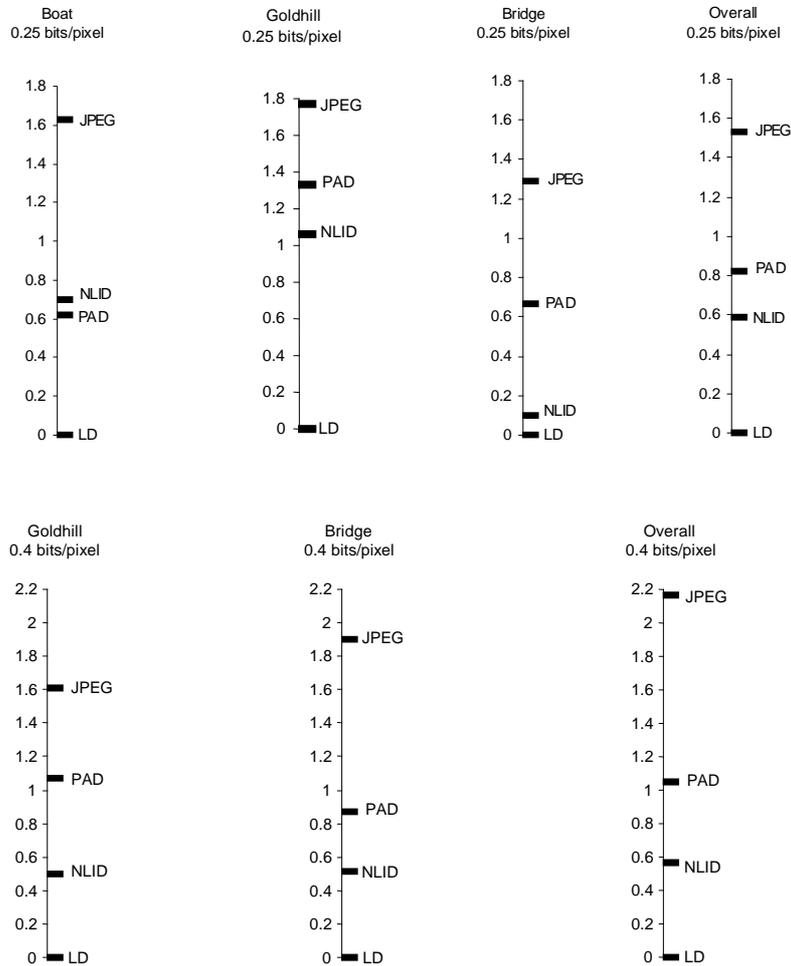


Figure 2 Scaling results. The scales show the expected values of a model observers quality assessment with respect to a chosen reference point, e.g., the position of LD. Positive difference on the scale means “better in expectance”. The unit is the uncertainty of the decision. Large differences mean a better discrimination.

subjective preference of the non-linear diffusion methods is significantly larger than that of the LD. The adaptive diffusions NLID and PAD are quite similar (compare Figs. 3 and 3) suggesting that the contrast adaptability is more important than the directional sensitivity, though since PAD is almost always better than NLID, directional adaptability is also important. We conclude that better adaptability leads to a better redistribution of the original bit/quality-rate by devoting larger portions of these resources on perceptually important information.

The compressed images *without* pre-processing show a superior subjective preference over the pre-processed ones, though they are actually worse in PSNR. This

Table 3 Results of the χ^2 -test of Mosteller for the subjective scaling experiment. The hypothesis is that the obtained scale values are correct. The sample size is n , the degree of freedom of the test is $df=3$. The probability that the scales are not random is $P(\chi^2)$. The regression coefficient $0 \leq r \leq 1$ gives the goodness of fit of the preference probabilities reconstructed from the model to those obtained by counting the scores.

		n	$P(\chi^2)$	r
0.25 bits/pixel	1. Boat	21	0.8611	0.9790
	2. Goldhill	21	0.72271	0.9627
	3. Bridge	21	0.79388	0.9627
	<i>Overall (1 & 2 & 3)</i>	<i>63</i>	<i>0.9985272</i>	<i>0.9965</i>
0.4 bits/pixel	1. Boat	21	0.033233	0.9665
	2. Goldhill	21	0.70797	0.9693
	3. Bridge	21	0.96986	0.9952
	<i>Overall (2 & 3)</i>	<i>42</i>	<i>0.99432</i>	<i>0.9986</i>

can be due to the fact that many subjects preferred sharp though falsely textured regions rather than smoothed areas, where the undersampled textures and compression-artifacts were removed by the pre-processing. It may follow the well-known effects of psycho visual illusions (like as Kanizsa figures [8]), which can be derived from the description of brain stimuli responses of neurosciences (e.g. [7]). It says that the brain, based on its high adaptability, may detect anomalous contours from the partially degraded details, like from artefacts.

Conclusion

We have considered the application of different diffusion methods as a pre-processing step for artefact reduction for the blockwise DCT compression, which we previously proposed in [21].

The compression quality improvement achievable by diffusion pre-processing is only up to 1.5% in PSNR for our test images and for low bit rates. For some of these test images, 2-3% improvement in PSNR was reported for post-processing methods [20] for the same bit rates. Note that diffusion pre-processing is done only once in the compression phase, and it consists of local iterative local operations. We have also seen that such small differences in the PSNR can yield surprisingly different results in the subjective evaluation.

The pre-processed images can be compressed with a significantly better PSNR than the original image, where the quality of all these compressed images have a better or equal value as compared to the original image. This means that these images are less susceptible to artefacts.

The subjective tests have shown that for a fixed bit rate and error-rate, adaptive diffusion captures better the visually important information in images. However,

since the diffusion pre-processing algorithm is tuned by using PSNR to estimate the image quality and the artefact reduction, and since the blur involved in the artefact-removal filtering is less tolerated by the human observers than the artefacts, the perceived subjective quality of the images will generally be decreased by pre-processing when done up to the maximal artefact reduction. We have to remark, that to our knowledge, no such subjective tests were done for other quality enhancement methods like post-processing.

Currently we are trying to find methods for establishing better diffusion pre-processing parameters in order to improve the subjective performance. We also do a modified subjective experiment, where the original image is displayed along with the two images to be compared. Preliminary results show that there are better parameter choices and that the two experiments, the one in the presence of the original image and the one without the original displayed, can yield quite different results.

References

1. L. Alvarez, F. Guichard, P. L. Lions, and J. M. Morel, "Axioms and fundamental equations of image processing", *Arch. Rational Mech. Anal.*, vol. 123, pp. 199-257, 1993.
2. L. Alvarez, P. L. Lions, and J. M. Morel, "Image selective smoothing and edge detection by non-linear diffusion II", *SIAM J. of Num. Anal.*, vol. 29, no. 3, pp. 845-866, 1992.
3. J. G. Apostolopoulos and N. S. Jayant, "Postprocessing for very low bit rate video compression", *IEEE Trans. Image Processing*, vol. 8, no.8, pp.1125-1129, 1999.
4. R. Castagno, S. Marsi, and G. Ramponi, "A simple algorithm for the reduction of blocking artifacts in images and its implementation", *IEEE Trans. on Consumer Electronics*, vol. 44, no.3, pp. 1062-1070, 1998.
5. F. Catté, T. Coll, P. L. Lions, and J. M. Morel, "Image selective smoothing and edge detection by non-linear diffusion", *SIAM J. Numerical Anal.*, vol. 29, pp.182-193, 1992.
6. L. Florack, "Image Structure", *Kluwer Academic Publishers*, 1997.
7. R. Heydt, E. Peterhans, "Mechanism of contours in monkey visual cortex. I. Lines of pattern discontinuity", *Journal of Neuroscience*, vol.9, pp.1731-1748, 1989.
8. G. Kanizsa, "Subjective Contours", *Scientific American*, vol.234, pp.48-52, 1976.
9. Y. L. Lee, H. C. Kim, and H. W. Park, "Blocking effect reduction of JPEG images by signal adaptive filtering", *IEEE Trans. on Image Processing*, vol.7, no.2, pp.229-234, 1998.
10. T. Lindeberg and B. M. Haar Romeny, "Linear scale-space I-II". In: *Geometry-Driven Diffusion In Computer Vision*, Kluwer Academic Publishers, pp.1-72, 1992.
11. J. Luo, C. W. Chen, K. J. Parker, and T. S. Huang, "Artifact Reduction in Low Bit rate DCT-Based Image Compression", *IEEE Trans. on Image Proc.*, vol.5, no.9, pp. 1363-1370, 1996.
12. J O. Monga, R. Deriche, G. Malandain and J.-P. Cocquerez, "Recursive filtering and edge tracking: two primary tools for 3-D edge detection", *Image and Vision Computing* vol. 4, no. 9, pp 203-214, 1991.
13. F. Mosteller: "Remarks on the method of paired comparisons III", *Psychometrika*, 16, pp. 207-218, 1951.
14. W. B. Pennebaker and J. L. Mitchell, "JPEG Still Image Data Compression Standard", *Van Nostrand Reinhold*, 1993.

15. M. K. Ozkan, M. I. Sezan, and A. M. Tekalp, "Adaptive motion-compensated filtering of noisy image sequences", *IEEE Trans. on Circuits and Systems for Video Tech.*, vol. 3, no. 4, pp. 277-290, 1993.
16. P. Perona and J. Malik, "Scale-space and edge detection using anisotropic diffusion", *IEEE Trans. Pattern Anal. Machine Intel.*, vol. 12, no. 7, pp. 629-639, 1990.
17. R. Prost and A. Baskurt, "JPEG dequantisation array for regularized decompression", *IEEE Trans. on Image Proc.*, vol. 6, no. 6, pp. 883-888, 1997.
18. T. Roska and T. Szirányi, "Classes of analogic CNN algorithms and their practical use in complex processing", In: *Proc. IEEE Non-linear Signal and Image Processing*, pp. 767-770, June, 1995.
19. K. Sauer, "Enhancement of low bit rate images using edge detection and estimation", *CVGIP: Graphical Models and Image Processing*, vol. 53, no. 1, pp. 52-65, 1991.
20. M. Y. Shen and C. C. J. Kuo, "Review of postprocessing techniques for compression artifact removal", *J. of Visual Comm. And Image Rep.*, vol. 9, no. 1, pp. 2-14, 1998.
21. T. Szirányi, I. Kopilović, and B. P. Tóth, "Anisotropic diffusion as a pre-processing step for efficient image compression", In: *Proc. ICPR*, pp. 1565-1567, Brisbane, Australia, August, 1998.
22. L. L. Thurstone, "A law of comparative judgment", *Psychol. Rev.*, 34, pp. 273-286, 1927.
23. W. S. Torgerson, "Theory and Methods of Scaling", *John Wiley and Sons, Inc.*, 1958.
24. Y. Yang, N. Galatsanos, and A. K. Katsaggelos, "Projection-based spatially adaptive reconstruction of block-transform compressed images", *IEEE Trans. on Circuits and Systems for Video Tech.*, vol. 3, no. 6, pp. 421-432, 1993.
25. Independent JPEG Group's JPEG, DJPEG, version 6a, 7-Feb-96, Copyright (C) 1996, Thomas G. Lane (<http://www.ijp.com>).



Figure 3 Maximal artifact reduction with pre-processing for “Goldhill” ($c=0.25$ bits/pixel, scale t_2). Details of the (a) JPEG compressed image, (b) pre-processing with LD, (c) NLID, (d) PAD.



(a)

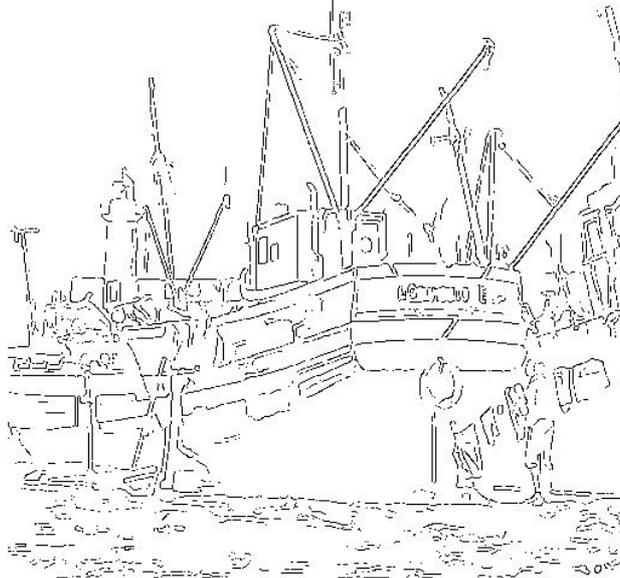


(b)

Figure 4 Boat at 0.25 bits/pixel. (a) JPEG, (b) JPEG with PAD preprocessing.

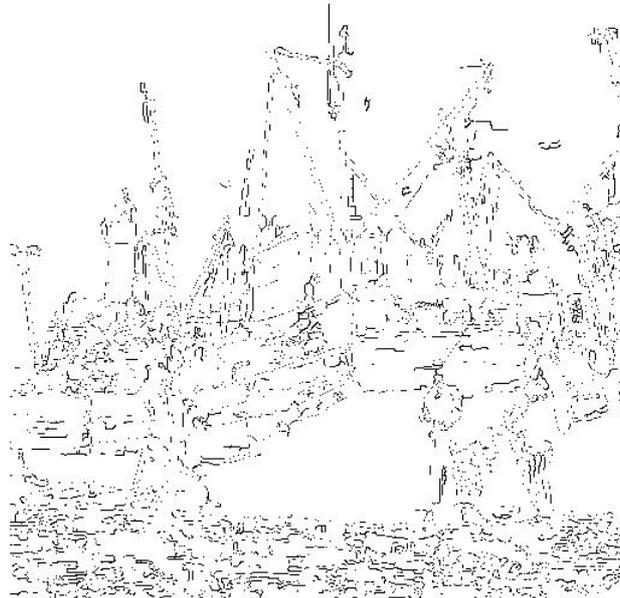


(a)

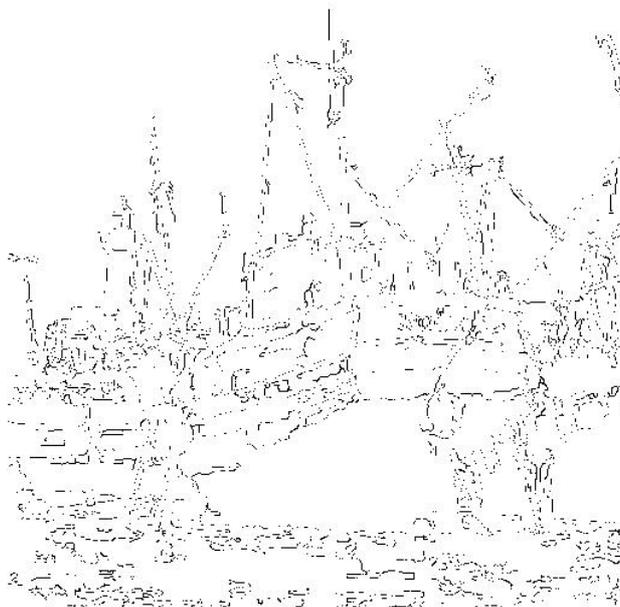


(b)

Figure 5 Edge detection for Boat 0.25 bits/pixel. (a) True edge points for JPEG, (b) true edge points for JPEG with PAD preprocessing.



(a)



(b)

Figure 6 Edge detection for Boat. (a) False edges for JPEG, (d) false edges for JPEG with PAD preprocessing.