EMPIRICAL EVALUATION OF BOUNDARY POLICIES FOR WAVELET–BASED IMAGE CODING

Claudia Schremmer
Praktische Informatik IV
Universität Mannheim, Germany
schremmer@informatik.uni-mannheim.de

Abstract - The wavelet transform has become the most interesting new algorithm for still image compression. Yet there are many parameters within a wavelet analysis and synthesis which govern the quality of a decoded image. In this paper, we discuss different image boundary policies and their implications for the decoded image. A pool of gray–scale images has been wavelet–transformed with different settings of the wavelet filter bank and quantization threshold and with three possible boundary policies.

Our empirical evaluation is based on three benchmarks: a first judgement regards the perceived quality of the decoded image. The compression rate is a second crucial factor. Finally, the best parameter settings with regard to these two factors is weighted with the cost of implementation.

Contrary to the new standard JPEG–2000, where mirror padding is implemented, our investigation proposes circular convolution as the boundary treatment.

Keywords: *Wavelet Analysis, Image Coding, Boundary Policies, Empirical Evaluation.*

1. INTRODUCTION

Due to its outstanding performance in compression, new image coding techniques such as the new standard JPEG–2000 [SCE00] [ITU00], focus on the wavelet transform (WT). As we were interested in the influence of the filter length on image coding quality — and contrary to the JPEG-2000 standard, where a reversible (Daubechies 5/3–tap) and an irreversible (Daubechies 9/7–tap) wavelet filter bank are proposed — we have investigated the orthogonal and separable wavelet filters developed by Daubechies [Dau92] which belong to the group of wavelets used most often in image coding applications. They specify a number n_0 of vanishing moments: if a wavelet has n_0 vanishing moments, then the approximation order of the wavelet transform is also n_0.

However, implementations of the WT on still images entail other aspects as well: speed, decomposition depth, and boundary treatment policies. Long filters require more computing time than short ones. Furthermore, the (dyadic) WT incorporates the aspect of iteration: the low–pass filter defines an approximation of the original signal which contains only half as many coefficients. This approximation successively builds the input for the next approximation. For compression purposes, coefficients in the time–scale domain are discarded and the synthesis quality improves with the number of iterations on the approximation. Finally, the WT is mathematically defined only *within* a signal; image applications thus need to solve the boundary problem. Depending on the selected boundary policy, the number of iterations in a WT might vary with the filter length. Moreover, the longer the filter length, the more important the boundary policy becomes.

In this work, we investigate the effects of three different boundary policies in combination with different wavelet filter banks on a number of gray–scale images. A first determining factor is the visual perception of a decoded image. As we will see, although the quality varies strongly with the selected image, for a given image it remains relatively unconcerned about the parameter settings. A second crucial factor is therefore the expected compression rate. Finally, the cost of implementation weights these two benchmarks. Our empirical evaluation leads us to recommend *circular convolution* as boundary treatment, contrary to JPEG–2000.

The article is organized as follows. In Section 2, we cite related work in wavelet filter evaluation. Section 3 reviews the wavelet transform and details the aspects that are important for our survey. In Section 4, we present the technical evaluation of the wavelet transform and detail our results. The article ends in Section 5 with an outlook on future work.

2. RELATED WORK

Villasenor's group researches wavelet filters for image compression. In [VBL95], the focus is on biorthogonal filters, and the evaluation is based on the information preserved in the reference signal, while [GFBV97] focuses on a mathematically optimal quantizer step size. In [AK99], the evaluation is based on lossless as well as on subjective lossy compression performance, complexity and memory usage. Interpretation of *why* the observations are made is nevertheless lacking. Strutz has thoroughly researched the dyadic WT in [Str97]: the design and construction of different wavelet filters is investigated, as are good Huffman and arithmetic encoding strategies. An investigation of boundary policies, however, is lacking. The new standard JPEG–2000 proposes mirror padding (or: periodic symmetric extension) as the image boundary treatment [SCE00] [ITU00].

3. THE WAVELET TRANSFORM

A wavelet is an (ideally) compact function, i.e., outside a certain interval it vanishes. Implementations are based on the fast wavelet transform, where a given wavelet (i.e., *mother wavelet*) is shifted and dilated so as to provide a base in the function space. In other words, a one–dimensional function is transformed into a two–dimensional space, where it is approximated by coefficients that depend on *time* (determined by the translation parameter) and on *scale*, i.e., frequency (determined by the dilation parameter). — By convention, the notion of time is used even for signals that depend on *location* rather than on time. Thus, a wavelet–transformed image is also said to be located in the *time*–scale domain. — The localization of a wavelet in time spread (σ_t) and frequency spread (σ_ω) has the property $\sigma_t \sigma_\omega = \text{const}$. However, the resolution in time and frequency depends on the frequency. This is the so–called *zoom* phenomenon of the WT: it offers high temporal localization for high frequencies while offering good frequency resolution for low frequencies. Consequently, the WT is especially well suited to analyze local variations such as those in still images: a high–frequency part of an image (e.g., a transition from bright foreground to black background) will be analyzed by short, high–amplitude wavelets. Low variations (e.g., gray value within the same object) will be analyzed by long, low–amplitude wavelets.

3.1 Wavelet Transform and Filter Banks

By introducing multiresolution, Mallat [Mal98] [Mal87] made an important contribution to the application of wavelet theory to multimedia: the transition from mathematical theory to filters. Multiresolution analysis is implemented via high–pass filters, resp. band–pass filters (i.e., wavelets) and low–pass filters (i.e., scaling functions). In this context, the wavelet transform of a signal can be realized with a filter bank via successive application of a 2–channel filter bank consisting of high–pass and low–pass filters: the detail coefficients (resulting from the application of the high–pass, resp. band–pass filter) of every iteration step are kept apart, and the iteration starts again with the remaining approximation coefficients (from application of the low–pass filter) of the transform. This multiresolution theory is 'per se' defined only for one–dimensional wavelets on one–dimensional signals. As still images are two–dimensional discrete signals and two–dimensional wavelet filter design

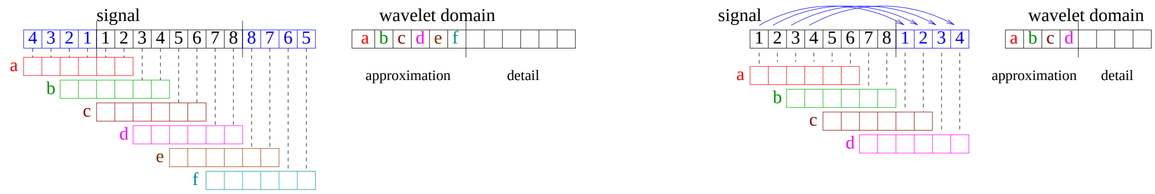

(a) Mirror Padding. (b) Circular Convolution.

Figure 1: Mirror padding versus circular convolution for a signal of length of 8 and a filter with 6 taps. Here, the filter is a low–pass filter, thus the coefficients resulting from the convolution form the approximation entries. In (a), the padding results in inflated storage space in the wavelet domain. In (b), the approximation contains half as many entries as the original signal. Together with the details, the entries of the wavelet domain require the same storage space as the original signal.

remains an active field of research [KS00] [KV95] [TK93], current implementations are restricted to *separable* filters. The successive convolution of filter and signal in both dimensions opens two potential iterations: standard decomposition (i.e., *all* approximations, even in mixed terms, are iterated) and non–standard decomposition (i.e., only the *purely* low–pass filtered parts of every approximation enter the iteration). In this work, we concentrate on the non–standard decomposition.

3.2 Image Boundary

A digital filter is applied to a signal by *convolution*. Convolution, however, is defined only *within* a signal. In order to result in a reversible wavelet transform, *each* signal coefficient must enter into **`filter_length/2`** calculations of convolution (here, the subsampling process by factor 2 is already incorporated). Consequently, every filter longer than two entries, i.e., every filter except *Haar*, requires a solution for the boundary. Furthermore, images are signals of a relatively short length (in rows and columns), thus the boundary treatment is even more important than e.g. in audio coding. Two common boundary policies are *padding* and *circular convolution*.

Padding Policies. With padding, the coefficients of the signal on either border are padded with **`filter_length-2`** coefficients (see Figure 1 (a)). Consequently, each signal coefficient enters into **`filter_length/2`** calculations of convolution, and the transform is reversible. Many padding policies exist: *constant padding*, where the signal's boundary coefficient is padded; *mirror padding*, where the signal is mirrored at the boundary; *spline padding*, where the border coefficients are extended by spline interpolation, etc. All padding policies have in common that each iteration step physically increases the storage space in the wavelet domain. In [Wic98], a theoretical solution for the required storage space (depending on the signal, the filter bank and the iteration level) is presented. Nevertheless, its implementation remains sophisticated.

Circular Convolution. The idea of circular convolution is to 'wrap' the end of a signal to its beginning or vice versa (see Figure 1 (b)). In so doing, circular convolution is the only boundary treatment to maintain the number of coefficients for a WT, thus simplifying storage management[1]. A minor drawback is that the time information contained in the time–scale domain of the wavelet–

[1] Storage space, however, expands indirectly: an image can be stored with `integers`, while the coefficients in the time–scale domain require `floats`.

transformed coefficients 'blurs': the coefficients in the time–scale domain that are next to the right border (resp. left border) also affect signal coefficients that are located on the left (resp. right).

The selected boundary policy has an important impact on the iteration behavior of the wavelet transform: convolving the signal with a filter is only reasonable for a signal length greater than the filter length, and each iteration step reduces the size of the approximating signal by a factor of 2. This does not affect the iteration behavior of padding policies. With circular convolution, however, the decomposition depth varies with the filter length: the longer the filter, the fewer decomposition iterations are possible. For example, for an image of 256×256 pixels, the Daubechies–2 filter bank with 4 taps allows a decomposition depth of 7, while the Daubechies–20 filter bank with 40 taps has reached signal length after only 3 decomposition levels, see the table below.

Filter Bank	Taps	It. levels
Daub–2	4	7
Daub–3	6	6
Daub–4	8	6
Daub–5	10	5
Daub–10	20	4
Daub–15	30	4
Daub–20	40	3

Thus, the evaluation presented in Tables 1 to 4 is based on a decomposition depth of level 8 for the two padding policies, while the decomposition depth for circular convolution varies from 7 to 3, according to the selected filter length.

4. EMPIRICAL EVALUATION

4.1 Setup

The goal of our empirical evaluation was to find the best parameter settings for the choice of the wavelet filter bank and for the image boundary policy to implement. The performance was evaluated according to the criteria:

1. visual quality,
2. compression rate,
3. complexity of implementation.

The quality was rated based on the peak signal–to–noise ratio (PSNR)[2]. The compression rate was simulated by a simple quantization threshold: the higher the threshold, the more coefficients in the time–scale domain are discarded, the higher is the compression rate. More precisely, the threshold was carried out only on the parts of the image that have been high–pass filtered (resp. band–pass filtered) at least once. In other words, the approximation of the image was excluded from the thresholding due to its importance for the image synthesis.

Our evaluation was set up on the six gray–scale images of size 256×256 pixels demonstrated in Figure 2. These test images have been chosen in order to comply with different features:

[2] When $\mathrm{org}(x, y)$ depicts the pixel value of the original image at position (x, y), and $\mathrm{dec}(x, y)$ denotes the pixel value of the decoded image at position (x, y), then

$$\text{PSNR [dB]} = 10 \cdot \log \left(\frac{\sum_{xy} 255^2}{\sum_{xy} (\mathrm{org}(x,y) - \mathrm{dec}(x,y))^2} \right).$$

- contain many small details: Mandrill, Goldhill,
- contain large uniform areas: Brain, Lena, Camera, House,
- be relatively symmetric at the left–right and top–bottom boundaries: Mandrill, Brain,
- be very asymmetric with regard to these boundaries: Lena, Goldhill, House,
- have sharp transitions between regions: Brain, Lena, Camera, House,
- contain large areas of texture: Mandrill, Lena, Goldhill, House.

4.2 Results

Image–Dependent Analysis.
The detailed evaluation results for the six test images are presented in Tables 1 and 2. Some interesting observations made from these two tables and their explanations are as follows:

- For a given image and a given quantization threshold, the PSNR remains astonishingly constant for different filter banks and different boundary policies.
- At high thresholds, 'Mandrill' and 'Goldhill' yield the worst quality. This is due to the large amount of details in both images.
- 'House' produces the overall best quality at a given threshold. This is due to its large uniform areas.
- Due to their symmetry, 'Mandrill' and 'Brain' show good quality results with padding policies.
- The percentage of discarded information at a given threshold is far higher for 'Brain' than for 'Mandrill'. This is due to the black uniform background of 'Brain', which produces small coefficients in the time–scale domain compared to the many small details in 'Mandrill' which produce large coefficients and thus do not fall below the threshold.
- With regard to the heuristic for compression, and for a given image and boundary policy, Table 2 reveals that
  - the compression ratio for zero padding *increases* with increasing filter length,
  - the compression ratio for mirror padding *decreases* with increasing filter length,
  - the compression ratio for circular convolution varies, but most often stays *almost constant*.

  The explanation is as follows. Padding an image with zeros, i.e., black pixel values, is most often a sharp contrast to the original image, thus the sharp transition between the signal and the padding coefficients results in large coefficients in the fine scales, while the coarse scales remain unaffected. This observation, however, is put into a different perspective for longer filters: with longer filters, the constant *run* of zeros at the boundary does not show strong variations, and the detail coefficients in the time–scale domain thus remain small. Hence, a given threshold cuts off fewer coefficients when the filter is longer. With mirror padding, the padded coefficients for shorter filters represent a good heuristic for the signal near the boundary. Increasing filter length and accordingly longer padded areas, however, introduces too much 'false' detail information into the signal, resulting in many large detail coefficients that 'survive' the threshold.

Image–Independent Analysis
The above examples reveal that most phenomena are signal–dependent. As a signal–dependent determination of best–suited parameters remains academic, our further reflections are made on the *average* image quality and the *average* amount of discarded information as presented in Tables 3 and 4 and the corresponding Figures 3 and 4.

Figure 3 visualizes the coding quality of the images, averaged over the six test images. The four plots represent the quantization thresholds 10, 20, 45 and 85. In each graphic, the visual quality (quantified via PSNR) is plotted against the filter length of the Daubechies wavelet filters. The three boundary policies: zero padding, mirror padding and circular convolution are regarded separately. The plots obviously reveal that the quality decreases with an increasing threshold. More important are the following statements:

- Within a given threshold, and for a given boundary policy, the PSNR remains almost constant. This means that the quality of the coding process does not or hardly depends on the selected wavelet filter bank.
- Within a given threshold, mirror padding produces the best results, followed by circular convolution. Zero padding performs worst.
- The gap between the performance of the boundary policies increases with an increasing threshold.

Nevertheless, the differences observed above with 0.28 dB maximum gap (at threshold = 85 and filter length = 40) are so marginal that they do not actually influence the visual perception.

As the visual perception is neither influenced by the choice of filter nor by the boundary policy, the coding performance has been studied as a second benchmark. The following observations are made from Figure 4. With a short filter length (4 to 10 taps), the compression ratio is almost identical for the different boundary policies. This is not astonishing, as short filters involve only little boundary treatment, and the relative importance of the boundary coefficients with regard to the signal coefficients is negligible. More important for our investigation are:

- The compression heuristic for each of the three boundary policies is inversely proportional to their quality performance. In other words, mirror padding discards the least number of coefficients at a given quantization threshold, while zero padding discards the most.
- With an increasing threshold, the gap between the compression ratios of the three policies narrows.

In the overall evaluation, we have seen that mirror padding performs best with regard to quality, while it performs worst with regard to compression. Inversely, zero padding performs best with regard to compression and worst with regard to quality. Circular convolution holds the midway in both aspects. On the other hand, the gap in compression is by far superior to the differences in quality. If we now call to mind the coding complexity of the padding approaches, compared to the ease of implementation of circular convolution (cf. Section 3.2), we strongly recommend to implement circular convolution as the boundary policy in image coding.

5. CONCLUSION

We have discussed and evaluated the strengths and weaknesses of different boundary policies in relation to different orthogonal wavelet filter banks. In opposition to the JPEG–2000 coding standard, where mirror padding is suggested for boundary treatment, we have proven that circular convolution performance is superior in the overall combination of quality performance, compression performance and ease of implementation.

In future work, we will improve our heuristic on the compression rate and rely on the calculation of a signal's entropy as it is presented in [WM01] and [Str97].

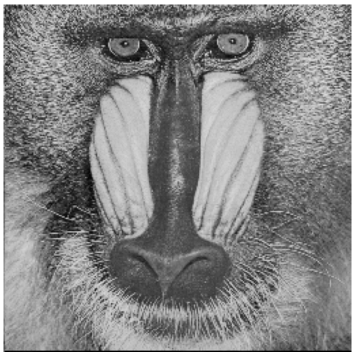

(a) Mandrill.

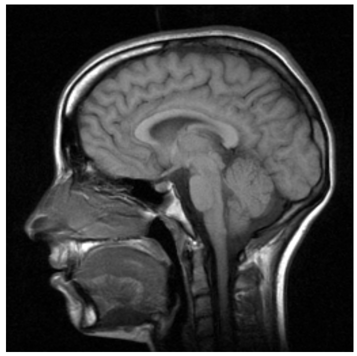

(b) Brain.

(c) Lena.

(d) Camera.

(e) Goldhill.

(f) House.

Figure 2: Test images. The test images (a) and (b) are relatively symmetric at their left–right and at their top–bottom boundary, while the images (c) — (f) are not.

Quality of visual perception — PSNR [dB]									
	Image								
	Mandrill			Brain			Lena		
Wavelet	zero padding	mirror padding	circular convol.	zero padding	mirror padding	circular convol.	zero padding	mirror padding	circular convol.
	Threshold: 10 — Excellent overall quality								
Daub–2	18.012	17.996	18.238	18.141	18.151	18.197	16.392	16.288	16.380
Daub–3	18.157	18.187	18.221	18.429	18.434	18.433	16.391	16.402	16.350
Daub–4	18.169	18.208	17.963	18.353	18.340	18.248	16.294	16.355	16.260
Daub–5	18.173	18.167	18.186	18.279	18.280	18.259	16.543	16.561	16.527
Daub–10	17.977	17.959	18.009	18.291	18.300	18.479	16.249	16.278	16.214
Daub–15	17.938	17.934	18.022	18.553	18.543	18.523	16.267	16.304	16.288
Daub–20	17.721	17.831	18.026	18.375	18.357	18.466	16.252	16.470	16.238
	Threshold: 20 — Good overall quality								
Daub–2	14.298	14.350	14.403	16.610	16.611	16.577	14.775	14.765	14.730
Daub–3	14.414	14.469	14.424	16.743	16.755	16.721	14.758	14.817	14.687
Daub–4	14.231	14.239	14.276	16.637	16.628	16.734	14.862	14.918	14.735
Daub–5	14.257	14.216	14.269	16.747	16.751	16.854	14.739	14.946	14.815
Daub–10	14.268	14.274	14.360	16.801	16.803	16.878	14.624	14.840	14.699
Daub–15	14.246	14.258	14.300	16.822	16.810	16.852	14.395	14.631	14.477
Daub–20	14.046	14.065	14.227	16.953	16.980	16.769	14.252	14.597	14.353
	Threshold: 45 — Medium overall quality								
Daub–2	10.905	10.885	10.910	14.815	14.816	14.747	13.010	13.052	12.832
Daub–3	10.988	10.970	10.948	15.187	15.150	15.052	12.766	13.138	12.903
Daub–4	10.845	10.839	10.885	15.014	15.029	15.056	12.820	13.132	12.818
Daub–5	10.918	10.969	10.949	15.036	15.031	14.999	12.913	13.301	12.983
Daub–10	10.907	10.929	10.913	14.989	15.013	15.212	12.447	13.066	12.795
Daub–15	10.845	10.819	10.815	15.093	15.133	15.064	12.577	12.954	12.686
Daub–20	10.784	10.872	10.843	14.975	14.934	14.882	12.299	12.877	12.640
	Threshold: 85 — Bad overall quality								
Daub–2	9.095	9.121	9.135	13.615	13.621	13.783	11.587	11.902	11.577
Daub–3	9.206	9.184	9.124	13.787	13.784	13.759	11.437	11.793	11.516
Daub–4	9.160	9.152	9.168	13.792	13.815	13.808	11.539	11.806	11.636
Daub–5	9.171	9.208	9.203	13.837	13.850	13.705	11.692	11.790	11.872
Daub–10	9.207	9.193	9.206	13.870	13.922	14.042	11.128	11.430	11.555
Daub–15	9.083	9.161	9.126	13.731	13.795	13.917	11.128	11.610	11.475
Daub–20	9.071	9.142	9.204	13.852	13.800	13.974	11.142	11.694	11.597
	Image								
	Camera			Goldhill			House		
Wavelet	zero padding	mirror padding	circular convol.	zero padding	mirror padding	circular convol.	zero padding	mirror padding	circular convol.
	Threshold: 10 — Excellent overall quality								
Daub–2	17.334	17.346	17.371	16.324	16.266	16.412	19.575	19.563	19.608
Daub–3	17.532	17.560	17.625	16.322	16.296	16.358	19.640	19.630	19.621
Daub–4	17.529	17.591	17.577	16.241	16.212	16.342	19.560	19.558	19.584
Daub–5	17.489	17.448	17.389	16.214	16.193	16.154	19.613	19.555	19.566
Daub–10	17.539	17.541	17.383	16.307	16.223	16.317	19.482	19.388	19.732
Daub–15	17.747	17.530	17.523	16.012	16.067	16.033	19.653	19.671	19.726
Daub–20	17.474	17.527	17.484	16.322	16.245	16.319	19.550	19.495	19.524
	Threshold: 20 — Good overall quality								
Daub–2	14.387	14.365	14.396	13.937	13.940	13.898	17.446	17.480	17.471
Daub–3	14.473	14.452	14.426	13.872	13.892	13.858	17.525	17.594	17.612
Daub–4	14.438	14.438	14.430	13.828	13.836	13.753	17.468	17.647	17.351
Daub–5	14.460	14.505	14.427	13.743	13.743	13.711	17.454	17.458	17.465
Daub–10	14.468	14.400	14.409	13.762	13.785	13.798	17.592	17.635	17.689
Daub–15	14.408	14.406	14.414	13.687	13.730	13.697	17.260	17.276	17.266
Daub–20	14.384	14.370	14.362	13.700	13.782	13.731	17.476	17.449	17.240
	Threshold: 45 — Medium overall quality								
Daub–2	12.213	12.242	12.131	12.033	12.034	11.876	15.365	15.437	15.155
Daub–3	12.032	12.122	12.188	11.961	12.006	11.889	14.957	15.476	15.118
Daub–4	12.150	12.178	12.145	11.855	11.891	11.925	14.906	15.080	15.180
Daub–5	12.077	12.133	12.120	11.848	11.844	11.801	15.159	15.382	15.244
Daub–10	12.061	12.197	12.093	11.760	11.917	11.726	14.776	15.246	14.872
Daub–15	12.074	12.059	12.176	11.725	11.855	11.753	14.810	15.090	14.969
Daub–20	11.798	11.975	12.048	11.763	11.803	11.703	14.420	15.033	14.609
	Threshold: 85 — Bad overall quality								
Daub–2	11.035	11.161	11.041	10.791	10.805	10.844	13.530	13.804	13.703
Daub–3	11.092	11.176	11.080	10.943	10.916	10.754	13.488	13.726	13.627
Daub–4	10.943	11.152	11.046	10.861	10.904	10.740	13.524	13.613	13.510
Daub–5	11.018	11.148	11.129	10.826	10.935	10.738	13.114	13.903	13.111
Daub–10	10.815	11.064	10.987	10.824	10.972	10.771	13.158	13.695	13.434
Daub–15	10.779	11.005	10.982	10.737	10.838	10.607	13.073	13.357	13.123
Daub–20	10.688	11.031	11.090	10.709	10.819	10.766	13.173	13.257	13.678

Table 1: Detailed results of the quality evaluation with the PSNR on the six test images. The mean values over the images for a fixed wavelet filter bank and a fixed boundary policy are given in Table 3.

Discarded information in the time–scale domain due to the threshold — Percentage [%]									
	Image								
	Mandrill			Brain			Lena		
Wavelet	zero padding	mirror padding	circular convol.	zero padding	mirror padding	circular convol.	zero padding	mirror padding	circular convol.
	Threshold: 10 — Excellent overall quality								
Daub–2	42	41	41	83	83	83	78	79	79
Daub–3	43	42	42	84	84	84	78	80	80
Daub–4	44	42	41	85	84	84	78	79	79
Daub–5	45	41	41	85	84	84	79	79	80
Daub–10	53	38	41	87	82	84	79	74	78
Daub–15	59	35	40	88	78	82	82	69	77
Daub–20	65	32	40	89	74	83	83	64	77
	Threshold: 20 — Good overall quality								
Daub–2	63	63	63	91	91	91	87	89	88
Daub–3	64	63	64	92	91	91	87	89	89
Daub–4	65	63	63	92	91	91	87	88	89
Daub–5	66	62	63	92	91	91	87	90	89
Daub–10	70	58	63	93	89	91	88	83	88
Daub–15	74	56	62	93	86	91	89	79	88
Daub–20	78	51	63	94	82	91	90	74	88
	Threshold: 45 — Medium overall quality								
Daub–2	86	86	87	96	96	96	94	95	95
Daub–3	86	86	87	96	96	96	94	95	95
Daub–4	87	86	87	96	96	96	94	95	96
Daub–5	87	85	87	96	96	96	95	94	96
Daub–10	88	82	87	97	94	96	94	91	96
Daub–15	90	79	87	97	91	96	95	88	96
Daub–20	92	74	87	97	89	96	96	83	96
	Threshold: 85 — Bad overall quality								
Daub–2	96	96	97	98	98	98	97	98	98
Daub–3	96	96	97	98	98	98	97	98	98
Daub–4	96	96	97	98	98	98	97	97	98
Daub–5	96	95	97	98	98	98	98	97	98
Daub–10	97	93	97	98	97	98	97	94	98
Daub–15	97	91	97	98	95	98	98	92	98
Daub–20	97	86	98	98	93	99	98	88	99
	Image								
	Camera			Goldhill			House		
Wavelet	zero padding	mirror padding	circular convol.	zero padding	mirror padding	circular convol.	zero padding	mirror padding	circular convol.
	Threshold: 10 — Excellent overall quality								
Daub–2	78	80	79	70	71	70	79	80	80
Daub–3	77	79	78	70	71	71	79	80	80
Daub–4	77	79	78	71	71	70	79	80	79
Daub–5	77	78	78	71	71	70	79	79	79
Daub–10	77	74	76	73	67	69	80	72	78
Daub–15	80	71	75	77	63	68	82	66	77
Daub–20	81	66	74	79	58	68	83	59	76
	Threshold: 20 — Good overall quality								
Daub–2	86	88	88	85	87	86	87	88	88
Daub–3	86	88	88	85	87	86	87	88	88
Daub–4	86	88	88	86	86	86	87	88	87
Daub–5	86	87	88	86	86	86	87	87	88
Daub–10	86	85	87	86	83	86	87	81	87
Daub–15	88	82	86	89	79	86	89	75	87
Daub–20	88	78	86	89	73	86	89	69	87
	Threshold: 45 — Medium overall quality								
Daub–2	93	95	95	94	96	95	93	95	94
Daub–3	93	95	95	95	96	95	94	95	95
Daub–4	94	95	95	95	95	95	94	94	95
Daub–5	94	94	95	95	95	96	94	94	95
Daub–10	93	93	95	95	92	96	94	89	95
Daub–15	94	91	95	95	89	96	95	84	94
Daub–20	95	88	95	96	85	96	95	78	95
	Threshold: 85 — Bad overall quality								
Daub–2	97	98	98	97	98	98	97	98	98
Daub–3	97	98	98	98	98	98	97	97	97
Daub–4	97	98	98	98	98	98	97	97	98
Daub–5	97	97	98	98	98	99	97	97	98
Daub–10	97	96	98	98	96	99	97	93	98
Daub–15	97	95	98	98	93	99	97	89	98
Daub–20	98	93	98	98	90	99	98	84	99

Table 2: Heuristic for the compression rate of the coding parameters of Table 1: the higher the percentage of discarded information in the time–scale domain is, the higher is the compression ratio. The mean values over the images for a fixed wavelet filter bank and a fixed boundary policy are given in Table 4.

Average image quality — PSNR [dB]						
Wavelet	zero padding	mirror padding	circular convol.	zero padding	mirror padding	circular convol.
	Threshold 10			Threshold 20		
Daub–2	17.630	17.602	17.701	15.242	15.252	15.246
Daub–3	17.745	17.752	17.768	15.298	15.330	15.288
Daub–4	17.691	17.711	17.662	15.244	15.284	15.213
Daub–5	17.719	17.701	17.680	15.233	15.270	15.257
Daub–10	17.641	17.615	17.689	15.253	15.290	15.306
Daub–15	17.695	17.675	17.686	15.136	15.185	15.168
Daub–20	17.616	17.654	17.676	15.135	15.207	15.114
	Threshold 45			Threshold 85		
Daub–2	13.057	13.078	12.942	11.609	11.736	11.681
Daub–3	12.982	13.144	13.016	11.659	11.763	11.643
Daub–4	12.932	13.025	13.002	11.637	11.740	11.651
Daub–5	12.992	13.110	13.016	11.610	11.806	11.626
Daub–10	12.823	13.061	12.935	11.500	11.713	11.666
Daub–15	12.854	12.985	12.911	11.422	11.628	11.538
Daub–20	12.673	12.916	12.788	11.439	11.624	11.718

Table 3: Average quality of the six test images. Figure 3 gives a more ‘readable’ plot of these digits.

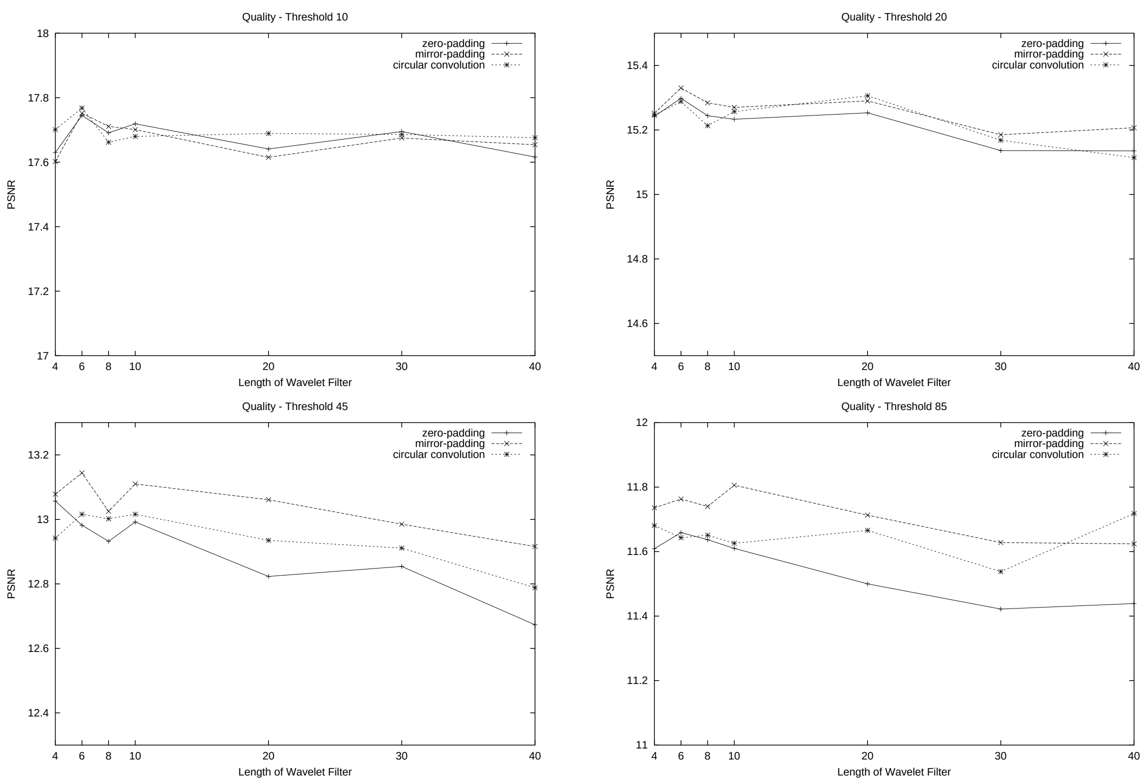

Figure 3: Visual quality of the test images at the quantization thresholds 10, 20, 45 and 85. The values are averaged over the six test images and correspond to Table 3.

Average discarded information — Percentage [%]						
Wavelet	zero padding	mirror padding	circular convol.	zero padding	mirror padding	circular convol.
	Threshold 10			Threshold 20		
Daub–2	72.0	72.3	72.0	83.2	84.3	84.0
Daub–3	71.8	72.7	72.5	83.5	84.3	84.3
Daub–4	72.3	72.5	71.8	83.8	84.0	84.0
Daub–5	72.7	72.0	72.0	84.0	83.8	84.2
Daub–10	74.8	67.8	71.0	85.0	79.8	83.7
Daub–15	78.0	63.7	69.8	87.0	76.2	83.3
Daub–20	80.0	58.8	69.7	88.0	71.2	83.5
	Threshold 45			Threshold 85		
Daub–2	92.7	93.8	93.7	97.0	97.7	97.8
Daub–3	93.0	93.8	93.8	97.2	97.5	97.7
Daub–4	93.3	93.5	94.0	97.2	97.3	97.8
Daub–5	93.5	93.0	94.2	97.3	97.0	98.0
Daub–10	93.5	90.2	94.2	97.3	94.8	98.0
Daub–15	94.3	87.0	94.0	97.5	92.5	98.0
Daub–20	95.2	82.8	94.2	97.8	89.0	98.7

Table 4: Average bitrate heuristic of the six test images. Figure 4 gives a more ‘readable’ plot of these digits.

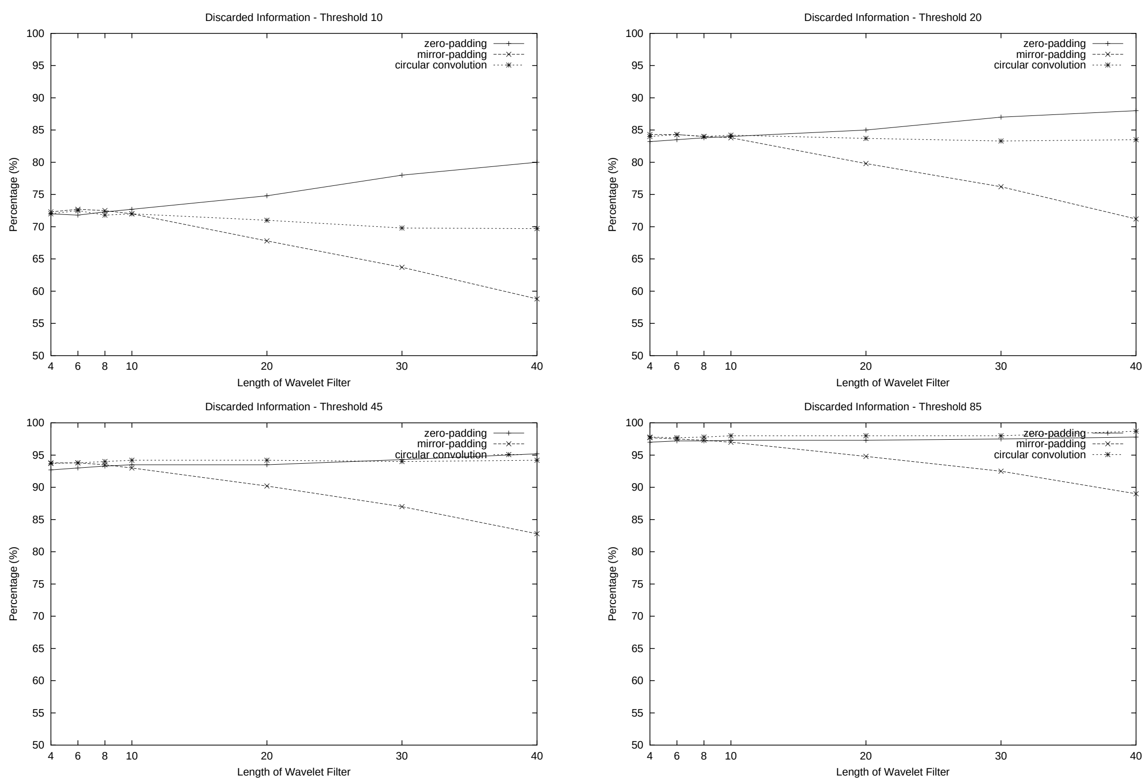

Figure 4: Average bitrate heuristic of the test images at the quantization thresholds 10, 20, 45 and 85. The values are averaged over the six test images and correspond to Table 4.

References

[AK99] Michael D. Adams and Faouzi Kossentini. Performance Evaluation of Reversible Integer–to–Integer Wavelet Transforms for Image Compression. In *Proc. IEEE Data Compression Conference*, page 514 ff, Snowbird, Utah, March 1999.

[Dau92] Ingrid Daubechies. *Ten Lectures on Wavelets*, volume 61. SIAM. Society for Industrial and Applied Mathematics, Philadelphia, Pennsylvania, 1992.

[GFBV97] Javier Garcia-Frias, Dan Benyamin, and John D. Villasenor. Rate Distortion Optimal Parameter Choice in a Wavelet Image Communication System. In *Proc. International Conference on Image Processing (ICIP)*, pages 25–28, Santa Barbara, CA, October 1997.

[ITU00] ITU. *JPEG 2000 Image Coding System*. International Telecommunication Union, Final Committee Draft Version 1.0 – FCD15444-1 edition, March 2000.

[KS00] Jelena Kovačević and Wim Sweldens. Wavelet Families of Increasing Order in Arbitrary Dimensions. *IEEE Transactions on Image Processing*, 9(3):480–496, March 2000.

[KV95] Jelena Kovačević and Martin Vetterli. Nonseparable Two– and Three–Dimensional Wavelets. *IEEE Transactions on Signal Processing*, 43(5):1269–1273, May 1995.

[Mal87] Stéphane Mallat. A Compact Multiresolution Representation: The Wavelet Model. *IEEE Computer Society Workshop on Computer Vision (WCV)*, 87:2–7, 1987.

[Mal98] Stéphane Mallat. *A Wavelet Tour of Signal Processing*. Academic Press, San Diego, CA, USA, 1998.

[SCE00] Athanassios N. Skodras, Charilaos A. Christopoulos, and Touradj Ebrahimi. JPEG2000: The Upcoming Still Image Compression Standard. In *11th Portuguese Conference on Pattern Recognition (RECPA00D)*, pages 359–366, Porto, Portugal, May 2000.

[Str97] Tilo Strutz. *Untersuchungen zur skalierbaren Kompression von Bildsequenzen bei niedrigen Bitraten unter Verwendung der dyadischen Wavelet–Transformation*. PhD thesis, Universität Rostock, Germany, May 1997.

[TK93] David B. H. Tay and N. G. Kingsbury. Flexible Design of Multidimensional Perfect Reconstruction FIR 2–Band Filters using Transformations of Variables. *IEEE Transactions on Image Processing*, 2(4):466–480, October 1993.

[VBL95] John D. Villasenor, Benjamin Belzer, and Judy Liao. Wavelet Filter Evaluation for Image Compression. *IEEE Transactions on Image Processing*, 2:1053–1060, August 1995.

[Wic98] Mladen Victor Wickerhauser. *Adapted Wavelet Analysis from Theory to Software*. A K Peters Ltd., Wellesley, Massachusetts, USA, 1998.

[WM01] Mathias Wien and Claudia Meyer. Adaptive Block Transform for Hybrid Video Coding. In *Proc. SPIE Visual Communications and Image Processing*, pages 153–162, San Jose, CA, January 2001.