WAVELET FILTER EVALUATION FOR IMAGE CODING

Claudia Schremmer, Christoph Kuhmünch, Christoph Esser
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: decomposition strategy, image boundary policy, quantization threshold, etc.

In this paper we discuss different image boundary policies and their implications for the decoded image. A focal point is the trade–off between the length of an orthogonal, compactly supported Daubechies–n wavelet filter bank and the decomposition depth of an image during analysis. An evaluation of the visual quality of images at different parameter settings leads to recommendations on the wavelet filter parameters to be used in image coding.

Keywords: *Wavelet Analysis, Decomposition Depth, Visual Quality, Evaluation.*

INTRODUCTION

Due to its outstanding performance in compression, new image coding techniques such as the new standard JPEG–2000 [1], focus on the wavelet transform (WT). The orthogonal and separable wavelet filters that Daubechies has developed 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. A fast approximation is mathematically desirable.

However, implementations of the WT on still images entail other aspects as well: speed, decomposition depth and boundary problems. Long filters require more computing time than short ones. Furthermore, the WT is mathematically only defined *within* a signal; image applications thus need to solve the boundary problem. Finally, the 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.

In this work, we investigate different wavelet filter banks in combination with different boundary policies. When *circular convolution* is chosen as the boundary treatment, the level of iteration depends on the length of the selected filter bank. We evaluate the trade–off between increasing coding quality by means of longer filters or increasing decomposition depth with shorter ones.

RELATED WORK

Villasenor's group researches wavelet filters for image compression. In [2], the focus is on biorthogonal filters, and the evaluation is based on the information preserved in the reference signal, while [3] focuses on a mathematically optimal quantizer step size. In [4], the evaluation is based on lossless as well as on subjective lossy compression performance, complexity and memory usage. Interpretation on *why* the observations are made is nevertheless lacking.

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 ('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 $\sigma_t \sigma_\omega =$ 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.

Wavelet Transform and Filter Banks

By introducing multiresolution, Mallat [5] 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 remains an active field of research [6], 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.

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 pixels of the signal on either border are padded with `filter_length-2` coefficients. 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 storage space in the wavelet domain is physically increased at each iteration step.

Circular Convolution. The idea of circular convolution is to 'wrap' the end of a signal to the beginning or vice versa. In so doing, circular convolution is the only boundary treatment to maintain the number of coefficients for a WT, thus simplifying storage management. However, the time information contained in the time–scale domain of the wavelet–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).

Iteration Behavior. 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.

EMPIRICAL EVALUATION

Our empirical evaluation was set up on a number of grayscale images of size 256×256 in order to find the best parameter settings on the choice of the wavelet filter bank and on the image boundary policy to implement. 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. The quality was rated based on the peak signal–to–noise ratio (PSNR)[1]. We have focused on the question whether the circular convolution with its ease of implementation and its drawback on the number of iteration levels provokes any loss of quality in the decoded image compared to padding.

As a thorough analysis of the results reveals that most phenomena are signal–dependent, we have decided *not* to average the results on all our test images, *but* to interpret the results based on the specific image. A general statement is possible, however. Table 1 shows the results of the three test images 'Baboon', 'Brain' and 'Lena' which are presented in Figure 1. The wavelet filter banks with the best results at a given parameter set of *image*, *boundary policy* and *threshold* are marked with *. The following observations are made from Table 1:

1. The PSNR for 'Brain' is generally higher than for 'Baboon' and 'Lena'.
2. 'Baboon' has a larger PSNR range than 'Lena': in good quality, the PSNR of 'Baboon' is higher, while 'Lena' wins this competition with regard to worse quality.
3. For 'Baboon' and 'Brain', the PSNR for circular convolution is generally slightly higher than for the other two policies; for 'Lena', mirror padding reaches the highest PSNR.
4. Most often, the PSNR is highest for a medium–length wavelet filter.

The explanations for these observations are as follows:

ad 1. The uniform black background of 'Brain' causes many coefficients in the time–scale domain to be small. This allows a high thresholding without deterioration.

ad 2. 'Baboon' contains a higher than average amount of details; they deteriorate the quality with high thresholding.

ad 3. The first two test images are (nearly) symmetric at the boundary, thus circular convolution is a suitable policy to concentrate the image's energy in a few high coefficients. Whereas 'Lena' is not symmetric.

ad 4. All wavelet filter banks have in common that the 'locality influence' of a coefficient in the time–scale domain corresponds to their length. The shorter filters are thus too irregular: they show strong block artifacts. The very long filters, however, 'blur' the locality information too much, intermixing different regions.

---

[1] 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
$\mathrm{PSNR} = 10 \cdot \log \left(\frac{\sum_{xy} 255^2}{\sum_{xy} (\mathrm{org}(x,y) - \mathrm{dec}(x,y))^2} \right)$.

Astonishingly enough, the number of possible iterations with circular convolution does not significantly influence the quality of the decoded image. Even with the non–symmetric 'Lena', the quality of circular convolution is sufficiently high. Its ease of implementation makes it especially suited to image coding. Contrarily to JPEG–2000, where mirror padding is proposed, the overall evaluation of quality and cost of implementation thus suggests to implement circular convolution.

Concerning the choice of wavelet filter, we recommend filters of medium length (10 to 20 taps), as their overall coding quality is superior to both shorter and longer filter banks.

CONCLUSION

We have discussed the strengths and weaknesses of different boundary policies in relation to different orthogonal wavelet filter banks. Circular convolution performs superior in the overall combination of ease of implementation and quality performance. The decreasing number of iterations that circular convolution implicates for increasing filter length can be disregarded in practical applications, as the recommendation is to implement orthogonal filters of medium length (Daubechies–5 to Daubechies–10).

References

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

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

[3] 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.

[4] 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.

[5] Stéphane Mallat. *A Wavelet Tour of Signal Processing.* Academic Press, San Diego, CA, USA, 1998. ISBN 0-12-466605-1.

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

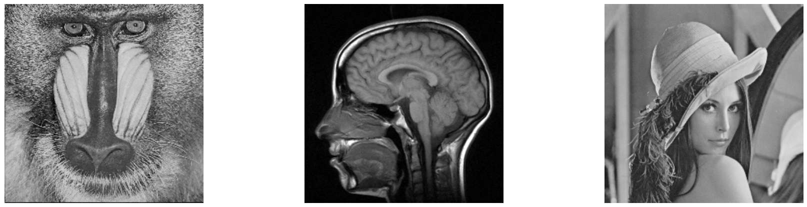

(a) Baboon. (b) Brain. (c) Lena.

Figure 1: Test images for Table 1.

Quality of visual perception — PSNR									
	Baboon			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

Table 1: Results of measurements of the images ‘Baboon’, ‘Brain’ and ‘Lena’. The PSNR (in dB) presents the quality of the visual perception.