



On the choice of ‘wavelet’ filters for still image compression

Olivier Rioul

► To cite this version:

Olivier Rioul. On the choice of ‘wavelet’ filters for still image compression. 1993 IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP-93), Apr 1993, Minneapolis, United States. pp.550-553, 10.1109/ICASSP.1993.319870 . hal-03330322

HAL Id: hal-03330322

<https://telecom-paris.hal.science/hal-03330322>

Submitted on 9 Aug 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L’archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d’enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

ON THE CHOICE OF “WAVELET” FILTERS FOR STILL IMAGE COMPRESSION

Olivier Rioul

France Télécom, CNET Paris B
CRPE (Unité mixte CNRS/CNET)
38-40 rue du Général Leclerc
92131 Issy-Les-Moulineaux, France.

ABSTRACT

“Regularity” is a new criterion brought by wavelet theory over filter banks. It is therefore important to know whether this criterion is relevant for applications such as image compression, in comparison with other filter properties. The following problem is addressed: How do regularity, frequency selectivity and phase act upon the performance of a still image compression scheme using wavelet decomposition? Preliminary results are given for a simple compression scheme using orthonormal separable wavelet transforms, scalar quantization, rate/distortion optimization, various coding criteria, and a large number of “wavelet” filters with balanced regularity, frequency selectivity and phase.

1. INTRODUCTION

This paper investigates the usefulness of several filter properties in a simple image compression scheme using a “discrete wavelet transform (DWT),” implemented as an octave-band tree filter bank allowing perfect reconstruction. Such a filter bank, depicted in Fig. 1, was successfully applied for some time in subband coding of speech and images. It is therefore clear that the DWT finds immediate application in compression problems: it is essentially a subband coding system. The main difference with traditional subband coding, and, therefore, the main novelty of wavelets in this context, is that filters are chosen to be “regular” [3,8].




Figure 1: Paraunitary two-band filter bank (in non-causal form) allowing perfect reconstruction. Filters $H(z)$ and $G(z)$ are half-band low-pass and high-pass, respectively.

Although intuitive arguments have been raised which hint that regularity should be useful for image coding, and several image compression schemes using regular wavelets have been proposed [1], the actual relevance of this new property is not clear yet [8]. It is therefore important to understand how this criterion acts upon the performance of

coding systems in competition with other constraints such as number of taps, frequency selectivity, and group delay deviation (filters whose phase is close to linear or not), and to measure its effects on compression performance.

2. REGULAR FILTER DESIGN

Today, there is a relatively small number of families of regular wavelet filters available in the literature—the most famous ones being Daubechies’ orthonormal filters [3]—and a number of compression schemes were designed using *ad-hoc* filters [1]. This is a serious limitation since different properties like frequency selectivity and regularity are interrelated inside one family of filters; hence, it is impossible to explain some coding performance as a consequence to one property and not to the other. We first overcome this limitation using a simple filter design procedure which allows one to vary stop-band attenuation, regularity, and phase quite independently.

We restrict ourselves to 1D FIR filters, hence use a separable wavelet transform, because regularity can be optimally measured in the 1D case [6], while estimating regularity is much more complex in the nonseparable case [2]. We also restrict ourselves to orthonormal filters, because only one filter has to be designed (analysis and synthesis filters are equal within time reversal). This brings better control over stop-band attenuation than in more general situations (“bi-orthogonal” filters [1], which are not half-band filters in general), but forbids (non-trivial) linear phase choices [3]. Note that it is also generally believed that the filter bank should not deviate far from orthonormality if efficient coding is needed [4]. This is anyway another open question.

Our design procedure maximizes stop-band attenuation for a given transition bandwidth, while imposing a zero of multiplicity K at half the sampling frequency in the low-pass frequency response. The latter condition ensures some regularity, which is quantified using an optimal algorithm proposed in [6]. Therefore, the obtained filters are maximally selective for a given regularity order.

The design constraints are linear in the coefficients of the product filter $P(z) = H(z)H(z^{-1})$ (regularity conditions become equality constraints, while frequency domain constraints become inequality constraints). Hence, the design problem can easily be solved using linear programming. More efficient design procedures are currently under study.

An extensive use of this procedure provides a large number of filters for which attenuation and regularity can be chosen quite independently with a good coverage. The design results can be summarized in figures such as Fig. 2 (a): For a given transition bandwidth, the curves provide the obtained attenuation vs. the regularity order. Notice that in Fig. 2 (a), regularity is quantified using Sobolev regularity, which is naturally related to magnitude specifications [7]. However, once the phase of the filters $G(z)$ and $H(z)$ are determined, regularity is best quantified using optimal Hölder regularity estimates [6,7]. The plots were, however, so close that we only provide the one making use of the Sobolev estimate.




Figure 2: (a). Attenuation (in dB) vs. “Sobolev” regularity [7] for the families of filters obtained. (b). Obtained low-pass filters frequency responses $|H(e^{j\omega})|^2$ with (normalized) transition bandwidth $\Delta\omega = 0.14$ and length $L = 16$. Daubechies and Smith-Barnwell filters (solid) correspond to $K = 8$ and 0 , respectively. Intermediate filters are shown for $K = 2$ (dashed), $K = 4$ (dotted) and $K = 6$ (dash-dotted).

Maximally selective regular filters of length $L = 16$ are shown in Fig. 2 (b). We observed that if $L/2 - K$ is odd, then the optimum filter provided by the design algorithm has one more zero at $z = -1$. For example, if one specifies $L = 16$ and $K = 1$, the optimum solution will have $K = 2$ zeroes at the Nyquist frequency. Therefore, only even values of K are considered in Fig. 2 (b). As K increases, stop-band attenuation is weaker, but flatness at $\omega = 0.5$ and regularity become higher. At one extreme, $K = L/2$, one recovers Daubechies filters [3]; at the other, $K = 0$, Smith

and Barnwell filters [9]. Therefore, this design procedure allows a soft transition between the two families of filters. Note that while Daubechies filters are not selective at all, selectivity can be greatly improved by relaxing a few zeroes at half the sampling frequency, resulting in a small loss of regularity. The converse statement is true for Smith and Barnwell filters.

Phase behavior: For a given frequency response, there are $2^{\lfloor L/4 \rfloor - 1}$ different filter solutions corresponding to different phases [3]. Unless otherwise mentioned, we use in this paper the solutions that are closest to linear phase (see e.g. [9]), i.e., whose group delay deviation in the pass-band is smallest (about 1 to 2.5 samples for $L \leq 16$). Therefore, there exists some flexibility concerning the choice of the phase. As shown in Fig. 3, solutions closest to linear phase are obtained for Daubechies filters.




Figure 3: Minimum group delay deviation (in the pass-band) of “closest to linear phase” filters $H(z)$ versus lengths L for several values of $\mu = L/4 - K/2$. Solutions closest to linear phase are obtained for $\mu = 0$ (Daubechies filters [3]).

3. RATE/DISTORTION OPTIMIZATION

The compression scheme (Fig. 4) consists of a separable DWT on J decomposition levels (or “octaves”), and a set of possible quantizers Q_i for each transformed subimage, corresponding to different bit rates (from 0 to 8 bits per pixel (bpp) with step size 0.5 bpp). The coder performance was evaluated by three different parameters, which can be taken as the bit rate reference R : overall quantizer bit rate, bit rate after Huffman coding and entropy. For simplicity, the results shown in this summary were obtained with scalar quantization and Huffman coding. However, we observed that results obtained with other coding criteria and/or lattice vector quantization in multidimensions [1] are similar except for the bit rate range.

In order to provide a fair comparison of compression results for different filters, we have used an optimization procedure similar to the one described in [5] for wavelet packets, which selects the best set of quantizers for each subimage and the best number of decomposition levels J which minimizes the overall distortion D at the reconstruction (measured by an m.s.e. criterion) for a given rate bud-




Figure 4: Image compression scheme with rate / distortion optimization.

get R_b . Thus the initial problem is

$$\min_{R \leq R_b} D.$$

The rate budget was chosen in the range $0.5 \leq R_b \leq 4$ bpp using the three bit rate definitions given above. This problem is solved by an unconstrained optimization procedure, using a Lagrangian cost function; it takes the form of a nested algorithm [5]

$$\max_{\lambda} \left\{ \min_{J, Q_i} D + \lambda R \right\} - \lambda R_b$$

where λ is the Lagrangian multiplier associated to the constraint $R \leq R_b$. As explained in [5], this algorithm is greatly simplified by the use of orthonormal filters, which make both rate and distortion additive over one step of decomposition: $R = \sum_i R_i$ and $D = \sum_i D_i$. We refer the interested reader to [5] for further details.

4. RESULTS

Fig. 5 shows a typical rate/distortion curve for the family of 12-tap filters designed as shown in section 2, and the 576×720 BARBARA image coded on 8 bpp. Clearly, trading regularity for selectivity in the filters affects the overall Peak SNR of the reconstructed image, for a wide range of bit rates: More regular (hence, less frequency selective) filters are best for a given number of taps. This was observed on various images, using various coding criteria, and relying either on the PSNR or on the visual quality of the reconstructed image. This is in agreement with the remark made by Kronander [4] that, surprisingly, “good” selectivity in frequency is not essential for coding performance, at least in the present framework of still image compression.

A detailed look at Fig. 5 reveals two categories of filters:

1. Those whose low-pass frequency response does not vanish at half the sampling frequency $\omega = 1/2$ ($K =$




Figure 5: PSNR vs. Huffman bit rate for 4 filters of length $L = 12$ and transition bandwidth $\Delta\omega = 0.0625$, corresponding to $K = 0, 2, 4$, and 6 (PSNR increases as K increases). Optimization was made using bit rate after Huffman coding as coding criteria.

0), resulting in very selective, but non-regular filters. The obtained PSNR curve lies below the ones corresponding to regular filters by more than 5dB. Finally, a strong artifact is clearly visible on the image, even for strongly attenuated filters (e.g., 40dB attenuation [9]) and was also observed by Kronander [4].

2. Regular filters ($K \geq 1$): Their performance all stand within only 1 to 2 dB difference for the same number of taps. It was nevertheless observed that the visual quality of the reconstructed image increases slightly as K increases (i.e., as frequency selectivity decreases). Here, the filters giving the best performance are Daubechies filters ($K = L/2$).

The following additional observations were also made for a number of different parameters in our compression scheme.




Figure 6: PSNR vs. Huffman bit rate for Daubechies filters of different lengths (L ranging from 2 to 18). PSNR first increases as length increases, until an asymptote is reached (for $L \approx 12$). Performance is not improving for longer filters (PSNR curves are even lowered a bit).

Fig. 6 illustrates the dependency of coding performance on filter lengths. For short filters, performance increases as length increases (which also increases regularity). However,

results are not improving greatly above $L = 10$ or 12 , which shows that using very regular filters is probably useless.




Figure 7: PSNR vs. Huffman bit rate for a “non-Daubechies” filter ($L = 12$, $K = 2$), designed with transition bandwidths $\Delta\omega = 0.0625$, 0.1 , and 0.14 . PSNR is globally higher for increasing values of $\Delta\omega$.

Fig. 7 compares coding performance for different transition bandwidths. We observed that performance is improved when increasing transition bandwidth $\Delta\omega$, provided that there are enough degrees of freedom in the filter design: Daubechies filters, for example, do not depend on the specified $\Delta\omega$. The dependency of performance on $\Delta\omega$ is, in fact, higher as the number of degrees of freedom, $L/2 - K$, is higher. Now, increasing $\Delta\omega$ improves stop-band attenuation, but also makes the frequency response of $H(z)$ closer to zero about the Nyquist frequency. This latter feature will generally increase regularity. Therefore, it seems that the relevant property in our framework is the ability of the low-pass filter frequency response to be very small about the Nyquist frequency, and a good measure for such behavior is regularity. Frequency selectivity is different, since it requires that the frequency response is attenuated over the entire stop-band, and it is generally believed that filters with “good” frequency selectivity should have a sharp transition bandwidth. The above discussion shows that frequency selectivity is not particularly useful here.

Finally, Fig. 8 compares coding performance for different phases, (where we have used the 256×256 LENA image for a change). The effect of phase on coding performance is almost unnoticeable for a fixed frequency response, either for the rate/distortion curves (less than 1 dB difference in Fig. 8) or for the visual quality of the reconstructed image. Although orthogonal filters cannot be linear phase for $L > 2$, one of the filters used to produce Fig. 8 has group delay variation less than 0.5 sample, which is fairly close to linear phase. Whether truly linear phase, biorthogonal filters can improve the situation remains an open question.

5. CONCLUSION

Results obtained for a simple compression scheme using various coding criteria, optimized rate/m.s.e. distortion, and a number of FIR filters with balanced regularity, frequency selectivity and phase, show that regularity may be relevant for still image compression, at least for short filters




Figure 8: PSNR vs. Huffman bit rate for the eight filter solutions of length $L = 14$ and $K = 5$, having the same magnitude response but different phase responses. The eight curves are almost undistinguishable at the level of the figure.

($L \leq 12$), for which the regularity order is relatively small. Using more regular filters is probably useless, as the compression performance is not improving greatly for longer filters. Moreover, the effect of phase seems negligible for orthonormal filters.

We should emphasize that our results, presented above, are valid only under the design assumptions and for the simple compression scheme described in the paper. The role of regularity for separable non-orthonormal linear phase systems and nonseparable systems is not investigated here and requires further investigation.

6. REFERENCES

- [1] M. Antonini, M. Barlaud, P. Mathieu, and I. Daubechies, “Image coding using vector quantization in the wavelet transform domain,” in *ICASSP'90*, Apr. 1990, pp. 2297–2300.
- [2] A. Cohen and I. Daubechies, “Non-separable bidimensional wavelet bases,” *Revista Matematica Iberoamericana*, 1992. To appear.
- [3] I. Daubechies, “Orthonormal bases of compactly supported wavelets,” *Comm. Pure Applied Math.*, vol. XLI, no. 7, pp. 909–996, 1988.
- [4] T. Kronander, *Some aspects of perception-based image coding*, PhD thesis, Linköping University, Sweden, 1989.
- [5] K. Ramchandran and M. Vetterli, “Best wavelet packet bases in a rate-distortion sense,” *IEEE Trans. on Image Proc.*, 1992. To appear.
- [6] O. Rioul, “A simple, optimal regularity estimate for wavelets,” in *Proc. European Signal Processing Conf. (EU-SIPCO)*, (Brussels), 1992.
- [7] O. Rioul, “Simple regularity criteria for subdivision schemes,” *SIAM J. Math. Anal.*, vol. 23, no. 6, pp. 1544–1576, Nov. 1992.
- [8] O. Rioul and M. Vetterli, “Wavelets and signal processing,” *IEEE Signal Processing Magazine*, vol. 8, no. 4, pp. 14–38, Oct. 1991.
- [9] M. J. T. Smith and T. P. Barnwell, “Exact reconstruction techniques for tree-structured subband coders,” *IEEE Trans. Acoust., Speech, Signal Processing*, vol. ASSP-34, no. 3, pp. 434–441, June 1986.