FOOTPRINTS AND EDGEPRINTS FOR IMAGE DENOISING AND COMPRESSION

Pier Luigi Dragotti[1] *Martin Vetterli*[1,2]

[1] Laboratoire de Communications Audiovisuelles
Swiss Federal Institute of Technology, CH-1015 Lausanne, Switzerland.
[2] EECS Dept., UC Berkeley CA 947020
e-mail: {dragotti,vetterli}@lcavsun1.epfl.ch
http://lcavwww.epfl.ch/~{dragotti,vetterli}/

ABSTRACT

In recent years wavelets have been quite successful in compression or denoising applications. To further improve the performance of wavelet based algorithms, we have recently introduced the notion of footprint, which is a data structure which contains all the wavelet coefficients generated by a discontinuity. The combined use of wavelets and footprints leads to very efficient algorithms for compression and denoising of 1-D piecewise smooth signals.

In this paper, we extend some of the previous results. We present a new denoising algorithm, where footprints are chosen adaptively according to the singularity locations. This new algorithm outperforms previously proposed ones. Then, we introduce the notion of edgeprints, which represents a natural extension of footprints to the two dimensional case. First experimental results on compression of 2-D piecewise smooth signal using edgeprints are promising.

1. INTRODUCTION

The design of a complete or overcomplete expansion that allows for compact representation of arbitrary signals is a central problem in signal processing applications. Wavelets, for instance, are known to be good approximants for 1-D piecewise smooth signals. The choice of a good basis, however, is only one of the elements that makes an efficient compression algorithm. In a recent work [3], we have introduced a compression algorithm that jointly uses wavelets and footprints. This algorithm has the right R-D behaviour for compression of piecewise polynomial signals. Footprint is a data structure that contains all the non-zero wavelet coefficients generated by a singularity.

It is of interest to compare footprints with other data structures like, for example, zerotrees [6]. Zerotrees indicate absence of singularities, while footprints describe singularities. Singularities are not well represented by zerotrees (instead, a number of non-zero coefficients are coded independently) while they can be more efficiently represented by footprints. [1]

In this work, we propose a different interpretation of footprints. We consider them as the elements of a redundant linear basis (frame). This expansion has some remarkable properties. It allows for a very compact representation of piecewise polynomial signals and it can be made locally orthogonal. This is important because, when compression or denoising are involved, it is more efficient to work with orthogonal bases. Thanks to these considerations, in Section 3, we present a new denoising algorithm which is adaptive in the choice of the footprint expansions. This new algorithm outperforms the non-adaptive one.

Finally, in Section 4, we introduce the notion of edgeprints which is a natural extension of footprints to the two dimensional case. This extension is similar to geometrical wavelets introduced in [5].

As footprints in 1-D, edgeprints allow to take advantage of the dependency between the wavelet coefficients generated by an edge in an image. This is an important condition to get an efficient 2-D compression algorithm.

2. THE FOOTPRINT EXPANSION

In [3], we have introduced the notion of footprints and proposed their use for compression or denoising of piecewise polynomial signals. What is interesting is that footprints can be seen as an overcomplete basis for representation of piecewise polynomial functions.

Consider a discrete time piecewise polynomial signal X with polynomials of maximum degree N. Let us perform a J level wavelet decomposition of X ($Y = WX$, where W is the wavelet operator). If the wavelet filter has at least $N+1$ vanishing moments then the discrete time poly-

Pier Luigi Dragotti was supported by Fonds National Suisse project number 20-061493.

[1] There are several other works on modeling of the correlation between wavelet coefficients. We do not quote all of them here for lack of space.

nomials are represented by the scaling coefficients, while the discontinuities are specified by few wavelet coefficients. Specifically: a time domain discontinuity generates a response of size J (the number of scales) by $L-1$ (L is the length of the wavelet filters) which we call a footprint. This footprint depends on the location of the discontinuity, namely, a particular discontinuity has 2^J footprints, depending on its location. From now on $f_k^{(N)}$ will denote a footprint; k is the discontinuity location in time domain and N the maximum degree of the two polynomials around the discontinuity.

It is interesting to notice that any footprint $f_k^{(N)}$ can be obtained as a linear combination of $N+1$ elementary orthonormal footprints: $\hat{f}_k^{(i)}$ $i=0,..,N$ ($\| \hat{f}_k^{(i)} \|=1$, and $< \hat{f}_k^{(i)}, \hat{f}_k^{(j)} >= \delta_{ij}$). That is: $f_k^{(N)} = \sum_{i=0}^{N} < Y, \hat{f}_k^{(i)} > \hat{f}_k^{(i)}$, where $< Y, \hat{f}_k^{(i)} >$ is the inner product between $\hat{f}_k^{(i)}$ and the wavelet coefficients of Y located at the same spatial position as the coefficients of $\hat{f}_k^{(i)}$. This is like saying that the wavelet coefficients generated by a time domain discontinuity lie on a $N+1$ dimensional sub-space and that the elementary footprints $\hat{f}_k^{(i)}$ are a possible basis for that subspace.

Now, if X has for instance only two discontinuities at positions k_1 and k_2, it clearly results, with a little abuse of notation:

$$X = W^{-1}\left(Y_s + \sum_{i=0}^{N}\alpha_{ik_1}\hat{f}_{k_1}^{(i)} + \sum_{i=0}^{N}\alpha_{ik_2}\hat{f}_{k_2}^{(i)}\right) \qquad (1)$$

where Y_s are the scaling coefficients of Y. This expression highlights the advantage of using the footprint expansion. A discrete time piecewise polynomial signal is completely characterized by the few footprint coefficients (and the scaling coefficients Y_s). Thus, footprints manage to give a sparse representation of piecewise polynomial signals.

If the two discontinuities are far enough then the corresponding footprints generated by them do not overlap and we say that they are orthogonal. Then $\alpha_{ik_1} =< Y, \hat{f}_{k_1}^{(i)} >$ and $\alpha_{ik_2} =< Y, \hat{f}_{k_2}^{(i)} >$ $i=0,..,N$. If the footprints overlap then they are not orthogonal and we need to construct the correct dual basis to find the right values for α_{ik_1} and α_{ik_2}.

Two footprints overlap if k_1 and k_2 are not far enough, but also if the number J of wavelet decomposition level is large. To reduce the probability of overlapping footprints J should be small. However, in some applications like signal denoising or compression performance increases if J is large. To overcome this trade-off we propose an adaptive choice of the level J according to the distance between discontinuities. A more detailed analysis of this adaptive algorithm is given in the next section, where a new denoising algorithm is presented.

3. ADAPTIVE FOOTPRINT EXPANSIONS FOR SIGNAL DENOISING

The term *denoising* is usually referred to the problem of removing noise from a corrupted signal without impairing the original signal too much. In the typical problem formulation the original signal X has been corrupted by additive Gaussian noise.

Now, assume that X is a piecewise polynomial signal defined over the interval $[0, T-1]$. We know that non-zero wavelet coefficients appear only around discontinuities. Moreover, if N is the maximum degree of any polynomial in the signal, then the wavelet coefficients generated by any discontinuity lie on a $N+1$ dimensional sub-space represented by the elementary footprints $\hat{f}_k^{(i)}$. Based on this consideration, in [4] it has been proposed to denoise the signal directly in the footprint domain. The advantage of doing so is that the dependency between wavelet coefficients is exploited.

Let us call F the noisy version of X. The non-adaptive denoising algorithm proposed in [4] works this way:

1. Compute a J level wavelet transform of the noisy signal: $G = WF$.
2. Define a threshold $T_h = \sigma\sqrt{2\ln T}$.
3. For each possible discontinuity position $k \in [1, T-1]$ compute the $N+1$ inner products $< G, \hat{f}_k^{(i)} >$, $i=0,..,N$.
4. Choose the location k_1 such that $\sum_{i=0}^{N} | < G, \hat{f}_{k_1}^{(i)} > |^2$ is maximum.
5. If
$$\exists i \in \{0..N\} : | < G, \hat{f}_{k_1}^{(i)} > | \geq T_h, \qquad (2)$$
then compute the residue:
$$R_G^1 = G - \sum_{i=0}^{D} < G, \hat{f}_{k_1}^{(i)} > \cdot \mathbf{1}\{| < G, \hat{f}_{k_1}^{(i)} | > T_h\}\hat{f}_{k_1}^{(i)}.$$
6. Iterate step 3-4-5 on the residue until condition (2) is not verified anymore.

Finally the estimated signal $\hat{X}$ is:

$$\hat{X} = W^{-1}\left[G_s + \sum_{m=0}^{M-1}\sum_{i=0}^{N} < R_G^m, \hat{f}_{k_m}^{(i)} > \hat{f}_{k_m}^{(i)}\right], \qquad (3)$$

where M is the total number of iterations, $R_G^0 = G$ and G_s are the scaling coefficients of G.

The performance of this algorithm depends on the choice of J. On one hand we want J to be large so that the estimation of the inner products $< G, \hat{f}_{k_1}^{(i)} >$ is more correct. [2] On the other hand if J is too large it can happen that the

[2] The estimation is improved with the size of the footprint vector.

wavelet coefficients generated by two discontinuities overlap and the proposed algorithm does not do the right estimation. In this case we need to know the dual basis of the overlapping footprints.

Now, assume that we have an estimation of the discontinuity locations of X. Call them $k_1, k_2, \ldots k_L$. In our adaptive algorithm, at each iteration, J is chosen to be the maximum possible value such that there are no overlapping footprints. That is: at the first iteration $J = J_1 = \lfloor \log_2(k_{p_1} - k_{p_1-1}) - \log_2(L-1) \rfloor$ where k_{p_1}, k_{p_1-1} are the two closest discontinuity locations. The non-adaptive algorithm is then applied on the interval $[k_{p_1-1}, k_{p_1}]$ with $J = J_1$ and the contributions of the estimated discontinuities are eliminated. The process is then iterated: the next chosen J is: $J_2 = \lfloor \log_2(k_{p_2} - k_{p_2-1}) - \log_2(L-1) \rfloor$ where k_{p_2}, k_{p_2-1} are the two closest discontinuity locations once the locations k_{p_1}, k_{p_1-1} have been eliminated. The non-adaptive algorithm is then applied on the new sub-interval $[k_{p_2-1}, k_{p_2}]$ and with $J = J_2$. The algorithm ends when all the discontinuities have been processed.

The last problem is to find a good estimator of the discontinuity locations $k_1, k_2, \ldots k_L$. In our approach we compute a $J = \log_2(T)$ level wavelet transform and construct the footprints expansion and the corresponding dual expansion. For each discontinuity position we compute the inner products between the wavelet coefficients and the right dual footprints. If for a given location k_p one of the $N+1$ inner products is larger than the threshold, we assume that at that position there is a time domain discontinuity.

In conclusion we can say that the estimation of the discontinuity locations is performed using the dual expansion, while the characteristic of the discontinuities (i.e. the value of the coefficients α_{ik_l} in Eq. (1)) is estimated using non-overlapping footprints.

4. EDGEPRINTS

Footprints are a powerful tool for signal compression and denoising, but they work only on one-dimensional signals. Thus, it is natural to look for an extension to the two dimensional case. The extension we propose is similar to the idea of bandelets introduced in [5].

Consider a two-dimensional piecewise polynomial signal with discontinuities represented by regular curves. [3] Each line (or each column) of this signal is piecewise polynomial, so it can be represented in terms of footprints. Consider the lth line and call it X, we have:

$$X = W^{-1}\left(Y_s + \sum_{i=0}^{N} \alpha_{l,k_1}^{(i)} \hat{f}_{k_1}^{(i)} + \sum_{i=0}^{N} \alpha_{l,k_2}^{(i)} \hat{f}_{k_2}^{(i)}\right). \quad (4)$$

[3] A real image can be approximately seen as two-dimensional piecewise polynomial signal, but where some pieces have added texture.

Where coefficients $\alpha_{l,k_1}^{(i)}$ for $i = 0, .., N$ characterize the singularity of the line l at position k_1. Now, the footprint coefficients detected are chained together to form curves. These curves which are one-dimensional signals are called *edgeprints*. Suppose that our image is formed by two polynomials of maximum degree N and these two polynomials are separated by one curve. This single curve generates $N+1$ edgeprints in much the same way that a singularity in the one-dimensional case generates $N+1$ footprint coefficients. The edgeprint of order zero is obtained by chaining the footprint coefficients of order zero (i.e. $\alpha_{x_1,x_2}^{(0)}$), the edgeprint of order one by chaining the footprint coefficients of order one and so on. Now, each edgeprint is a one-dimensional piecewise smooth function that can be efficiently compressed using footprints. Thus, applying footprints on edgeprints we get a very sparse representation of the original image.

The main advantage of using edgeprints instead of bandelets is that we completely exploit the dependency between the wavelet coefficients generated by an edge. In fact, the wavelet coefficients generated by an edge are spatially correlated and correlated across scales. By applying footprints on edgeprints, we fully exploit either the spatial correlation or the correlation across scales.

Our compression algorithm works this way:

1. Use footprints on each line and on each column of the image to detect singularities. Any singularity is characterized by $N+1$ footprint coefficients.
2. Chain the footprint coefficients to form curves. The chaining of coefficients of the same order gives an edgeprint. Each curve generates $N+1$ edgeprints.
3. Treat each edgeprints as a 1-d signal and compress it using footprints.
4. Compress the information related to the geometry of the curve using a sort of vertex-based shape coding [7]

The decoder first reconstruct each edgeprints, then locate each of them in the right position using the geometrical information.

5. SIMULATION RESULTS

Denoising: The denoising algorithm presented in Section 3 is based on critically subsampled footprints. That is, on footprints built on critically subsampled wavelet coefficients. We have also performed an extension of this denoising algorithm to the case of non-subsampled footprints. For lack of space we have not presented this extension here, but we will include this system in our simulation results.

In Table 1 we compare the performance of our denoising systems against a classical hard thresholding algorithm [2] and against cycle-spinning [1]. In this experiment we consider piecewise linear signals with no more than three discontinuities. The performance is analyzed in function of the size T of the signal. The table clearly shows that subsampled footprints system outperforms hard thresholding system, while non-subsampled footprints outperform cycle-spinning. In Fig 1, we show an example of the denoising algorithm.

T	64	128	256	512
Subsampled Footprints	16.2dB	18.5dB	19.8dB	22.1dB
Hard Thresholding	12.9dB	15.2dB	16.6dB	19dB
Cycle Spinning	16.3dB	18.6dB	20.3dB	22.9dB
Non-Subsampled Footprints	17.1dB	19.6dB	21dB	23.2dB

Table 1. Piecewise linear signals with four different polynomials.

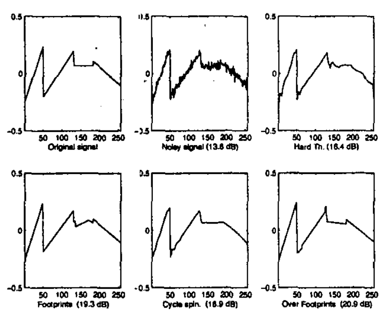

Fig. 1. SNR results. a) Original Signal. b) Noisy Signal (13.8dB). c) Hard Th. (16.4dB). d) Subsampled Footprints (19.3dB). e) Cycle Spin. (18.9dB). f) Non-Subsampled Footprints (20.9dB).

Compression: To show the potentialities of edgeprints, we show this simple experiment. Consider the image shown in Fig. 2(a). It is composed of three constant regions separated by two smooth curves. In one of the regions we have added "texture", which is simply represented by additive noise. The edgeprints generated by the two curves are approximately piecewise constant and constant. The coder compress the two edgeprints using footprints and then sends, as geometric information, the coordinates of few points of the curves. The decoder reconstruct the curves approximating them with piecewise linear functions. In this case the coder approximates the smooth curve with six points and the straight curve with two points. For the texture, the coder estimates the variance of the noise and send a quantized version to the receiver. The decoder generates locally the noise with the quantized version of the variance. In Fig. 2(b) we show the approximation of the original image obtained using edgeprints (before compression) and in Fig. 2(c) the error image. As expected the edgeprints well approximate the smooth behaviour of the signal. The error image just contains texture. The effect of compression is finally shown in Fig. 2(d).

The problem of finding the right allocation of bits to code the geometrical information, the edgeprints and the texture is a topic under investigation.

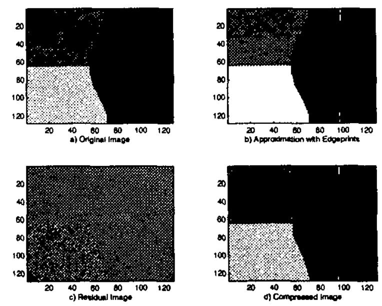

Fig. 2. Compression using Edgeprint: a) Original image, b) Approximation with Edgeprints, c) Residual Image, d) Compressed Image (bit budget=154, SNR=25.1dB).

6. CONCLUSIONS

In this paper we have proposed a new denoising algorithm based on footprints. The main innovation of this algorithm is that footprints are chosen according to the estimated distance between singularities. Moreover, we have investigated a possible generalization of footprints to the 2-D case. First experimental results are promising.

7. REFERENCES

[1] R.R. Coifman and D.L. Donoho "Translation-invariant denoising," Wavelets and Statistics, Lecture Notes in Statistics. SpringerVerlag, 1995.

[2] D.L.Donoho and I.M. Johnstone "Ideal spatial adaptation via wavelet shrinkage," Biometrika, 81:425-455, 1994.

[3] P.L. Dragotti and M. Vetterli "Wavelet transform footprints: catching singularities for compression and denoising", Int. Conf. on Image Processing, Vancouver, September 2000.

[4] P.L. Dragotti and M. Vetterli "Shift-Invariant Gibbs Free Denoising Algorithm based on Wavelet Transform Footprints", In Proc. SPIE2000, Wavelet Application in Signal and Image Processing, August 2000.

[5] E. Le Pennec and S. Mallat "Image Compression with Geometrical Wavelets", Int. Conference on Image Processing, Vancouver, September 2000.

[6] J.M. Shapiro, "Embedded image coding using zerotrees of wavelets coefficients," IEEE Trans. Signal Processing, vol. 41 pp.3445-3462, Dec. 1993.

[7] G.M. Schuster, Gerry Melnikov, and Aggelos K. Katsaggelos "Operational Optimal Vertex-Based Shape Coding", IEEE Signal Proc. Magazine, 91-108, November 1998.