Directional Wavelets and Wavelet Footprints for Compression and Denoising

Pier Luigi Dragotti[1], Martin Vetterli[1,2], Vladan Velisavljevic[1]

[1] Laboratoire de Communications Audiovisuelles
Swiss Federal Institute of Technology,
CH-1015 Lausanne, Switzerland.
Voice: +41 21 6937663, Fax: +41 21 6934312
[2] EECS Dept., UC Berkeley CA 947020

E-mail: *{dragotti,vetterli,velisavljevic}@lcavsun1.epfl.ch*
Web: *http://lcavwww.epfl.ch/*

ABSTRACT

In recent years, wavelet based algorithms have been successful in different signal processing tasks. The wavelet transform is a powerful tool because it manages to represent both transient and stationary behaviours of a signal with few transform coefficients. In this paper we present new expansions and algorithms which improve wavelet algorithms. First we focus on one dimensional piecewise smooth signals and propose a new representation of these signals in terms of elements which we call footprints. Then we consider two dimensional signals and present a new directional wavelet transform, which keeps the simplicity of the standard separable wavelet transform but allows for more directionalities. Denoising and compression algorithms based on footprints and directional wavelets show interesting improvement over traditional wavelet methods.

1 INTRODUCTION

The design of a complete or overcomplete expansion that allows for compact representation of certain relevant classes of signals is a central problem in signal/image processing and its applications. Parsimonious representation of data is important for compression [1], furthermore, achieving a compact representation of a signal also means intimate knowledge of the signal features and this can be useful for many other tasks including denoising, classification and interpolation.

The success of the wavelet transform is mainly due to its ability to characterize certain classes of signals with few transform coefficients. In particular, wavelets as time frequency localized bases are particularly suited to represent piecewise smooth functions and this is in contrast with the Fourier bases, which are inadequate when discontinuities are present. Although wavelets represent piecewise smooth signals well, the wavelet coefficients generated by discontinuities are dependent across scales. Thus, it is possible to characterize them more efficiently.

We present a new way to represent piecewise smooth signals in terms of objects which we call *footprints* and which make up an overcomplete dictionary of elements. The footprints dictionary is built from the wavelet transform. Given a signal of interest, we first perform the wavelet transform of this signal and then the wavelet coefficients are expressed in terms of footprints. The main property of footprints is that they characterize efficiently the singular structures of the signal, which usually carry important information. By constructing the footprint expansion on the wavelet transform, we remove the dependency across scales of the wavelet coefficients completely. Thus, by representing any discontinuity with the combination of a few footprints, we can get a sparser representation of the signal under consideration.

In two dimensions (2-D), the situation is much more open. In fact, wavelets are not good at modelling piecewise smooth signals (where discontinuities are along smooth curves). The two dimensional wavelet transform is a separable transform given by the tensor-product of two one dimensional (1-D) wavelets along the horizontal and vertical direction. For this reason, this separable transform is good at isolating horizontal and vertical edges, but it is not adequate at treating more complex discontinuities.

Non-separable approaches [2], in particular using directional filter banks [3, 4], have been investigated, showing the potential of truly non-separable methods. Such methods come at a price in terms of design and computational complexity. Some separable approached have been made in [5] but not on discrete spaces.

In our work, we wish to retain the simplicity of the separable wavelet transform while realizing some of the potential of non-separable schemes. We do this by introducing a directional wavelet transform that acts much like a standard separable transform but allows more directionalities. This is done by introducing *digital directions* that partition the discrete plane. Along these directions, it is then possible to apply the wavelet transform or the footprint expansion. Many useful properties are again inherited from the one-dimensional case.

The paper is organized as follows. The next section is dedicated to one dimensional signals. We introduce the footprint expansion and develop footprint based algorithms for compression and denoising of piecewise smooth signals. Section 3 focuses on two-dimensional signals. We define digital directions and the associated partitions of $\mathbb{Z}^2$. We also show how directional elements in an image are treated by separable directional transforms. In Section 4, we present numerical results for both one dimensional and two dimensional signals. We conclude in Section 5.

2 Representation of 1-D signals

We start by studying one dimensional signals and then we move to the two-dimensional case. In the next section, we demonstrate a decomposition of a piecewise smooth signal into a piecewise polynomial signal and a regular residual (Theorem 1). We then introduce the notion of footprints and present footprint based algorithms for denoising and compression.

2.1 Signal models

We consider 1-D piecewise smooth signals, that is, signals which are made of smooth pieces. For example, we define a piecewise smooth function $f(t)$, $t \in [0,T]$ with $K+1$ pieces, as follows

$$f(t) = \sum_{i=0}^{K} f_i(t)\mathbf{1}_{[t_i,t_{i+1}]}(t), \tag{1}$$

where $t_0 = 0$, $t_{K+1} = T$ and $f_i(t)$ is uniformly Lipschitz α over $[0,T]$ [1]. Those signals are interesting, because many signals encountered in practice can be well modeled as piecewise smooth. There is also another set of functions we will consider and which form the more restricted class of piecewise polynomial signals. A function $p(t)$ $t \in [0,T]$ is piecewise polynomial with $K+1$ pieces if

$$p(t) = \sum_{i=0}^{K} p_i(t)\mathbf{1}_{[t_i,t_{i+1}]}(t), \tag{2}$$

where $t_0 = 0$, $t_{K+1} = T$ and $p_i(t)$, $\ i = 0,1,..K$ are polynomials of maximum degree D.

Despite their simplicity, piecewise polynomial signals represent an important tool to characterize the nonstationary behaviour of piecewise smooth functions. It follows [7]

Theorem 1 *Given is a piecewise smooth signal $f(t)$ defined as in Eq. (1), that is, with pieces of Lipschitz regularity α. Then, there exists a piecewise polynomial signal $p(t)$ with pieces of maximum degree $d = \lfloor \alpha \rfloor$ such that the difference signal $r_\alpha(t) = f(t) - p(t)$ is uniformly Lipschitz α over $[0,T]$.*

Theorem 1 indicates a practical way to deal with piecewise smooth signals. It shows that any piecewise smooth signal $f(t)$ can be expressed as the sum of a piecewise polynomial signal and a residual which is uniformly Lipschitz α. That is

$$f(t) = p(t) + r_\alpha(t).$$

Now, since the residual is regular, it can be well represented with wavelets (the wavelet decomposition of $r_\alpha(t)$ results in small coefficients with fast decay across scales). Therefore, the only elements we need to analyse are discontinuities in the piecewise polynomial function and, in particular, the dependency across scales of the wavelet coefficients generated by these piecewise polynomial discontinuities.[2] In the next section, we present a new way to express discontinuities in piecewise polynomial signals. Together, with Theorem 1, this will lead to efficient algorithms to represent piecewise smooth signals. Although, we could perform this analysis in continuous time, we concentrate on the discrete-time case. This is because our final target is to develop efficient algorithms that act on discrete-time signals.

2.2 Wavelet Footprints: Theory

We move from continuous-time to discrete-time signals and introduce the notion of footprints which are scale-space vectors containing all the wavelet coefficients generated by particular polynomial discontinuities.[3] We show that any piecewise polynomial discontinuity is specified by the linear combination of a few footprints, and that footprints can be interpreted as an overcomplete expansion with good approximation properties.

Consider, first, a piecewise constant signal $x[n]$, $n \in [0, N-1]$ with only one discontinuity at position k. Consider a J level wavelet decomposition of this signal with a Haar wavelet:

$$x[n] = \sum_{j=1}^{J}\sum_{l=0}^{N/2^j-1} y_{jl}\psi_{jl}[n] + \sum_{l=0}^{N/2^J-1} c_l\phi_{Jl}[n]. \tag{3}$$

where $y_{jl} = \langle x, \psi_{jl}\rangle$, and $c_l = \langle x, \phi_{Jl}\rangle$ [4]. Now, since the Haar wavelet has one vanishing moment and finite support, the non-zero wavelet coefficients of this decomposition are only in the cone of influence of k. Thus Eq. (3) becomes

$$x[n] = \sum_{j=1}^{J} y_{jk_j}\psi_{jk_j}[n] + \sum_{l=0}^{N/2^J-1} c_l\phi_{Jl}[n],$$

where $k_j = \lfloor k/2^j \rfloor - 1$. Moreover, all these coefficients depend only on the amplitude of the discontinuity at k. Thus, if one defines a vector which contains all of them, one can specify any other step discontinuity at k by multiplying this vector by the right factor. This consideration leads to the following definition:

Definition 1 *Given a piecewise constant signal x with only one discontinuity at position k, we call footprint $f_k^{(0)}$ the norm one scale-space vector obtained by gathering together all the wavelet coefficients in the cone of influence of k and then imposing $\|f_k^{(0)}\| = 1$.*

Expressed in the wavelet basis, this footprint can be written as $f_k^{(0)}[n] = \sum_{j=1}^{J} d_{jk_j}\psi_{jk_j}[n]$, where $d_{jk_j} = y_{jk_j}/\sqrt{\sum_{j=1}^{J} y_{jk_j}^2}$. Now, any piecewise constant signal $x[n]$ with a step discontinuity at k can be represented in

[1] For a definition of Lipschitz regularity refer to [6].

[2] For simplicity, we call *piecewise polynomial discontinuity* a singularity between two polynomials.

[3] In continuous time, one can define footprints equivalently, but they are of infinite dimension and so of little computational value.

[4] Note that we are assuming N to be a power of 2.

terms of the scaling functions $\phi_{Jl}[n]$ and of $f_k^{(0)}$. For instance, the signal $x[n]$ in Eq. (3) becomes

$$x[n] = \sum_{l=0}^{N/2^J-1} c_l \phi_{Jl}[n] + \alpha f_k^{(0)}[n],$$

where $\alpha = \langle x, f_k^{(0)} \rangle = \sum_{j=0}^{J} y_{jk_j} d_{jk_j}$. That is, $\langle x, f_k^{(0)} \rangle$ represents the inner product between $f_k^{(0)}$ and the wavelet coefficients of x located at the same scale-space positions of the coefficients of $f_k^{(0)}$. The above discussion can be repeated for any other step discontinuity at different locations. For each location l we have a different footprint $f_l^{(0)}$. Now, given the complete dictionary $\mathcal{D} = \{f_k^{(0)}, k = 1, 2, ..N-1.\}$ of footprints, we can express any piecewise constant signal in terms of the elements of this dictionary and of the scaling functions.

The notion of footprints can be easily generalized to the case of piecewise polynomial signals (for more details refer to [7]). In this case, it can be shown that the wavelet coefficients in the cone of influence of a polynomial discontinuity at location k have only $D+1$ degrees of freedom (D is the maximum degree of any polynomial in the signal). Thus, one can characterize this discontinuity using $D+1$ footprints $f_k^{(d)}$, $d = 0, 1, .., D$. Moreover, these footprints are designed in a way to guarantee the two following conditions

$$\begin{aligned} &\|f_k^{(d)}\| = 1 \quad d = 0, 1, .., D; \\ &\langle f_k^{(i)}, f_k^{(j)} \rangle = \delta_{ij} \quad i = 0, 1, .., D; \;\; j = 0, 1, .., D. \end{aligned} \tag{4}$$

To characterize any polynomial discontinuity, we need a dictionary $\mathcal{D} = \{f_k^{(d)}, d = 0, 1, .., D; k = 0, 1, .., N-1.\}$ of $(D+1)N$ footprints. With this dictionary of footprints and with the scaling functions, we can represent any piecewise polynomial signal. In particular, a piecewise polynomial signal x with K discontinuities at locations $k_1, k_2, .., k_K$ is given by

$$x[n] = \sum_{l=0}^{N/2^J-1} c_l \phi_{Jl}[n] + \sum_{i=0}^{K} \sum_{d=1}^{D} \alpha_i^{(d)} f_{k_i}^{(d)}[n]. \tag{5}$$

Footprints are orthogonal to the scaling functions. Footprints related to the same locations are orthogonal too (see Eq. (4)). But footprints related to close discontinuities are biorthogonal. In particular, we have that $\langle f_l^{(d)}, f_k^{(c)} \rangle = 0$ for $|l-k| > (L-1) \cdot 2^J$, where L is the length of the wavelet filter. Thus, the orthogonality of footprints depend on the number J of wavelet decomposition level. Now, assume that we know the discontinuity locations and call k_{m-1}, k_m the two closest discontinuities in (5). If J is chosen such that $J = \lfloor \log_2(k_m - k_{m-1}) - \log_2(L-1) \rfloor$, than we are sure that footprints related to locations $k_1, k_2, .., k_K$ are orthogonal. In the next section, we present an iterative denoising algorithm where the number J is chosen adaptively according to the distance between discontinuities. In this way, there are no biorthogonal footprints to represent that signal. It is also of interest to note that footprints manage to provide a sparser representation of piecewise polynomial signals than the wavelet transform. Moreover, when $J = \log_2 N$, footprints form an unconditional basis for piecewise polynomial signals. That is, any linear combination of footprints gives a signal which is piecewise polynomial [7].

2.3 Wavelet Footprints: Applications

We focus on two main applications for which wavelets are successful, namely denoising and compression. We present alternative algorithms based on the footprint expansion and show that these methods can further improve wavelet based algorithms. The main characteristic of the footprint methods is that they can deal more efficiently with discontinuities.

2.3.1 Denoising

The term denoising usually refers to the removal of noise from a corrupted signal. In the typical problem formulation, the original signal x has been corrupted by additive noise. One observes $y[n] = x[n] + e[n]$ where $e[n]$ are independent and identically distributed (i.i.d.) zero mean Gaussian variables with variance σ^2 and the original signal is deterministic and independent of the noise. The goal of the denoising algorithm is to obtain an estimate $\hat{x}$ of the original signal which minimizes a risk function, usually the mean square error $E[\| x - \hat{x} \|^2]$. The wavelet based denoising algorithm introduced by Donoho and Johnstone [8] simply shrinks the wavelet coefficients. That is, it sets all wavelet coefficients smaller than a threshold to zero and keeps the coefficients above the threshold (hard thresholding) or shrinks them by a fixed amount (soft thresholding). The threshold is usually set to $T = \sigma\sqrt{2 \ln N}$, where N is the size of the signal [8]. A limit of this approach is that it does not exploit the dependency across scales of the wavelet coefficients. Thus, to overcome this limit, we apply a threshold in the footprint domain rather than in the wavelet domain. Doing so, we better exploit the dependency of the wavelet coefficients across scales. As a matter of fact, denoising in the footprint domain is equivalent to applying a vector threshold in the wavelet domain rather than a scalar threshold as in the usual methods.

Assume that $x[n]$ is piecewise polynomial. We can express piecewise polynomial signals in terms of footprints, thus our denoising system attempts to estimate this footprint representation from the observed noisy version $y[n]$.

For simplicity let us focus on piecewise constant signals. We first perform an estimation of the discontinuity locations and then we estimate the values of the footprints coefficients $\alpha_i^{(0)}$. The discontinuity locations are estimated in the following way:

1. Choose a dictionary $\mathcal{D} = \{f_k^{(0)} = \sum_{j=1}^{J} d_{jk_j} \psi_{jk_j}; k = 0, 1, ..N-1\}$ of footprints with $J = \log_2 N$. This dictionary represents a biorthogonal basis.

2. Compute the dual basis of $\mathcal{D}$ and call $\tilde{f}_k^{(0)}$ $k = 1, 2, .., N-1$ the elements of this dual basis [5].

[5] It is of interest to emphasize that this dual basis turns out to be a

3. Compute the $N-1$ inner products $\langle y, \tilde{f}_k^{(0)} \rangle$ $k = 1, 2, .., N-1$.

4. Consider as discontinuity locations the ones related to the inner products larger than the threshold $T_k = \|\tilde{f}_k^{(0)}\| T$. That is, if $|\langle y, \tilde{f}_k^{(0)} \rangle| \geq T_k$, then assume that there is a discontinuity at location k. T is the universal threshold ($T = \sigma\sqrt{(2\ln N)}$) [8]

Now, we have a set of estimated discontinuity locations: $\hat{k}_1, \hat{k}_2, .., \hat{k}_{\hat{K}}$. The problem is that, due to the noise, this estimation can have errors. Thus, this problem must be considered in the next step where the footprints coefficients are evaluated.

1. Given the set of estimated discontinuity locations, take $J_1 = \lfloor \log_2(\hat{k}_m - \hat{k}_{m-1}) - \log_2(L-1) \rfloor$, where $\hat{k}_{m-1}, \hat{k}_m$ are the two closest estimated discontinuity locations.

2. For each possible location $k \in [\hat{k}_{m-1}, \hat{k}_m]$ compute the inner product $\langle y, \frac{\hat{f}_k^{(0)}}{\|\hat{f}_k^{(0)}\|} \rangle$, where $\hat{f}_k^{(0)}$ is the sub-footprint obtained by considering only the first J_1 wavelet coefficients of $f_k^{(0)}$. That is: $\hat{f}_k^{(0)} = \sum_{j=1}^{J_1} d_{jk_j} \psi_{jk_j}[n]$

3. Choose the location k_1 such that $|\langle y, \frac{\hat{f}_{k_1}^{(0)}}{\|\hat{f}_{k_1}^{(0)}\|} \rangle|$ is maximum.

4. If
$$|\langle y, \frac{\hat{f}_{k_1}^{(0)}}{\|\hat{f}_{k_1}^{(0)}\|} \rangle| \geq T, \quad (6)$$
then compute the residue:
$$R_y^1 = y - \frac{1}{\|\hat{f}_{k_1}^{(0)}\|} \langle y, \frac{\hat{f}_{k_1}^{(0)}}{\|\hat{f}_{k_1}^{(0)}\|} \rangle f_{k_1}^{(0)}.$$

5. Iterate step 3-4 on the residue until condition (6) is not verified anymore.

6. Once condition (6) is not verified anymore, remove the two discontinuity locations $\hat{k}_{j-1}, \hat{k}_j$. If the set of remaining discontinuity locations is not empty, find a new decomposition level J_2 and go to step 2. Otherwise, if all discontinuities have been considered, stop.

Finally, the estimated signal $\hat{x}$ is:
$$\hat{x} = < y, \phi_J > \phi_J[n] + \sum_{m=0}^{M-1} \frac{1}{\|\hat{f}_{k_m}^{(0)}\|} < R_y^m, \frac{\hat{f}_{k_m}^{(0)}}{\|\hat{f}_{k_m}^{(0)}\|} > f_{k_m}^{(0)}[n], \quad (7)$$
where M is the total number of iterations and $R_y^0 = y$.

First, notice that, since the footprints $f_{k_m}^{(0)}$ in Eq. (7) are obtained taking a wavelet transform with $J = \log_2 N$ decomposition level, then we are sure that the estimated signal $\hat{x}$ is piecewise constant as x. This is an important property, because in this way we are sure that the estimated signal does not present artifact around discontinuities (pseudo-Gibbs effect). This algorithm can be easily generalized to the case of piecewise polynomial signals and, thus, we do not detail this generalization here.

Now, assume that the original signal $x[n]$ is piecewise smooth. In this case, we use a two step denoising algorithm. First, we estimate the piecewise polynomial behaviour of x using this footprints based algorithm. Then we use the standard thresholding method based on the wavelet transform to denoise the residual $r[n] = y[n] - \hat{x}[n]$.

Denoising in the wavelet domain suffers from the lack of shift invariance of the wavelet basis. One way to overcome this limitation is to use a denoising method called cycle-spinning [9]. For a range of shifts, cycle spinning shifts the noisy signal, denoises each shifted version and, finally, unshift and average the denoised signals. Since footprints suffer from the same lack of shift invariance as wavelets, one can use the idea of cycle spinning to reduce this shift dependency. The only difference between cycle spinning with wavelets and cycle spinning with footprints is that, in this second case, each shifted version of the signal is denoised with footprints rather than wavelets.

In Section 4, we consider both methods (denoising with footprints and cycle spinning with footprints) and compare them with the equivalent wavelet based algorithms.

2.3.2 Compression

Wavelets are widely used in compression. The reason is that wavelets have very good approximation properties for representing certain classes of signals like piecewise smooth signals. While good approximation properties are necessary for good compression, it might not be enough. In compression, one has to consider the costs corresponding to indexing and compressing the retained elements in the approximation and independent coding of these coefficients might be inefficient [10].

Consider a piecewise smooth signal defined as in Eq. (1), that is, a function with pieces that are α-Lipschitz regular and with a finite number of discontinuities. It was shown in [11] that standard wavelet based schemes such as zerotrees [12] can achieve the following distortion-rate performance:
$$D(R) \leq c_1 R_s^{-2\alpha} + c_2 \sqrt{R_e} 2^{-c_3\sqrt{R_e}}, \quad (8)$$
where $R = R_s + R_e$ and R_e are the bits used to quantize the wavelet coefficients generated by the discontinuities, while R_s are the bits used to code the wavelet coefficients corresponding to the smooth parts of the signal. Now, suppose that the signal is piecewise polynomial. Then the wavelet coefficients related to the smooth parts of the signal are exactly zero, and so there is no need to use any bits to code them. The distortion of a wavelet based scheme becomes
$$D(R) \leq c_2 \sqrt{R} 2^{-c_3\sqrt{R}}. \quad (9)$$
However, a direct approach to compression of piecewise polynomial signals, based on an oracle telling us where

first order derivative.

discontinuities are, will lead to $D(R) \leq c_4 2^{-c_5 R}$ [13] and such behaviour is achievable using dynamic programming [13]. This large gap between ideal performance given by the scheme based on dynamic programming and wavelet performance is mainly due to the independent coding of the wavelet coefficients across scales. Statistical modeling [14] of such dependencies can improve the constants in (9), but going from $\sqrt{R}$ to R in the exponent requires taking the deterministic behaviour of wavelet coefficients across scales at singularities into account. This is well done using footprints, which thus close the gap with the ideal performance [7]

Theorem 2 *Consider piecewise polynomial signals with polynomials of maximum degree D and no more than K discontinuities. A coder, which represents these signals in the footprints basis and which scalar quantizes the discontinuity locations and the footprint coefficients achieves*

$$D(R) \leq c_6 2^{-c_7 R}. \qquad (10)$$

Thus, this theorem shows that, in case of piecewise polynomial signals, footprints significantly improve performance of wavelet coders. Footprints can be used for piecewise smooth signals too. Theorem 1 shows that a piecewise smooth signal can be separated into two contributions a piecewise polynomial part $p[n]$ and a residual $r[n]$ which is regular (α-Lipschitz over $\mathbb{R}$). Now, $p[n]$ can be compressed with footprints and this coder achieves (10). The residual $r[n]$ can be compressed with any other coder which achieves [11]

$$D(R) \leq c_8 R^{-2\alpha}. \qquad (11)$$

It is worth noticing that, because of the regularity of $r[n]$, the performance in (11) can be achieved with a simple coder based on linear approximation of $r[n]$ in a wavelet or Fourier basis [11]. Combining (10) and (11) shows that a two stage compression algorithm based on footprints and on linear approximation of the residual achieves

$$D(R) \leq c_8 R_s^{-2\alpha} + c_6 2^{-c_7 R_e}. \qquad (12)$$

Comparing (8) and (12), we can seen that this coder does not change the asymptotics of the distortion-rate function of wavelet coders ($\sim c_8 R_s^{-2\alpha}$). But, by coding the discontinuities efficiently, this coder reaches the asymptotic behaviour more rapidly. Finally, notice that, for this last performance, the underlying assumption is that the encoder knows in advance the signal to code, in this way it can separate the polynomial and the smooth parts of the signal. In the experimental results, we will show that a realistic encoder can obtain similar performance without knowing the signal characteristics in advance.

3 EXTENSIONS TO 2-D SIGNALS

In this section, we briefly outline, possible extensions of the previous results to the case of 2-D signals. Recall that our target is to keep the simplicity of separable transforms while realizing some of the potential of non-separable transforms. We define digital directions and associated partitions of $\mathbb{Z}^2$ and also show, with an illustrative example, how directional elements in an image are treated by a separable directional transforms.

A line in $\mathbb{R}^2$ is a simple object, but its equivalent in $\mathbb{Z}^2$ is a bit more complicated, as well known in raster graphics for example. We define a digital line of angle θ as a one-dimensional set of pixels approximately along a line of angle θ. In addition, the digital line and its shifts along orthogonal direction have to tile $\mathbb{Z}^2$. While there are many solutions to this problem, a simple and tractable one is to use the analytical definition of a discrete line [15]. The line is determined by its slope and shift by the following equation:

$$y[n] = \lfloor kx[n] \rfloor + \lfloor B \rfloor \qquad (13)$$

where $k = tan(\theta)$ represents the slope and belongs to the range $0 \leq k \leq 1$, and B represents the real-valued shift parameter. The definition of a discrete approximation of a line insures that each pixel belongs exactly to one line for a chosen slope. Lines with slopes out of the range may be obtained by symmetry, rotating and flipping vertically the space. This gives access to a wealth of directions in $\mathbb{Z}^2$.

Now, the standard separable wavelet transform simply consists of two transforms along the horizontal and vertical directions. This only permits the characterization and compression of phenomena along those two directions.

In our approach, one is free to choose any of the digital directions given by Eq. (13) and to apply the wavelet transform (or any other one dimensional transform such as footprints) along that direction. Moreover, the process can be iterated and at each iteration one can choose a new direction. Thus, building multi-resolution decompositions along multiple directions permits characterization and compression of phenomena other than just horizontal and vertical ones, as in the standard separable wavelet transform.

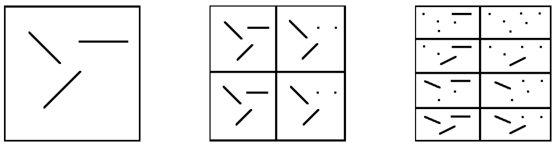

a) *b)* *c)*

Figure 1: A simple object and its standard and directional transform; (*a*) Original image, (*b*) standard, horizontal-vertical transform, 1 step, (*c*) 3-direction transform, 1 step.

Figure 1 conceptually shows this. Figure 1(b) shows the one level wavelet decomposition of the image given in Figure 1(a) with the standard separable wavelet transform. In this case we have four channels and only horizontal and vertical directions are efficiently isolated. Figure 1(c) depicts a one level wavelet decomposition of the same image but with wavelets transforms taken along three different directions. This second approach does isolate the different directions, as well as combinations thereof, and this in an intuitive way.

4 SIMULATION RESULTS

4.1 Denoising with footprints

In this case, we consider only piecewise polynomial signals. In Table 4.1, we compare the performance of our denoising systems with a classical hard thresholding algorithm [8] and cycle-spinning [9]. In this experiment, we consider piecewise linear signals with no more than three discontinuities. The performance is analysed in function of the size N of the signal. The table clearly shows that denoising with footprints outperforms the hard thresholding system, while cycle-spinning with footprints outperform traditional cycle-spinning. In Figure 2, we show an example of the denoising algorithm on a piecewise linear signal. We can see that signals denoised with footprints present better visual quality since they do not suffer from pseudo-Gibbs effects.

N	64	128	256	512
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
Cycle-spin footprints	17.1dB	19.6dB	21dB	23.6dB

Table 4.1. Denoising of piecewise linear signals with no more than three discontinuities.

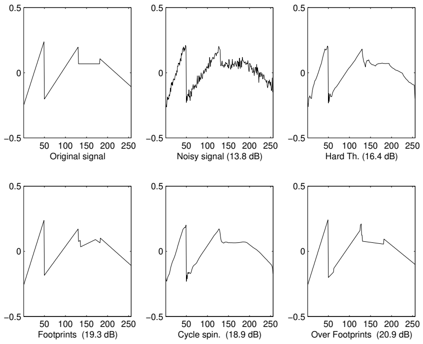

Figure 2: SNR results. a) Original Signal. b) Noisy Signal (13.8dB). c) Hard Th. (16.4dB). d) Subsampled Footprints (19.3dB). e) Cycle Spin. (18.9dB). f) Non-Subsampled Footprints (20.9dB).

4.2 Compression with footprints

In Theorem 2, we have shown that in case of piecewise polynomial signals, a footprint based coder can achieve the ideal rate-distortion performance. That is, it has the correct rate of decay of the R-D function. Now, we are interested in a numerical confirmation of this theorem. We consider piecewise constant signals with no more than five discontinuities. The signal has size $N = 2^{10}$ and the discontinuity locations are uniformly distributed over the interval $[0, N-1]$. The footprint coder operates as in Theorem 2, that is, it scalar quantizes the footprint coefficients and the discontinuity locations. Bits are allocated with a reverse waterfilling strategy. In Figure 3, we compare the rate-distortion performance of this footprint coder against the ideal bound and the ideal performance of a wavelet based coder. We can see that the behaviour of the footprint coder is consistent with the theory, since it has the same rate of decay as the ideal distortion function. Finally, we consider a piecewise smooth signal. The compression operates in the following way. With a denoising-like algorithm, we estimate the piecewise polynomial behaviour underlying the signal and compress it with footprints. The residual is assumed regular and it is compressed in a wavelet basis. That is,

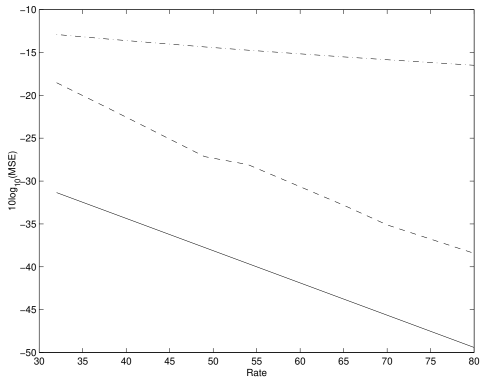

Figure 3: Theoretical and experimental D/R curves. Dashed-dotted: theoretical wavelet performance, dashed: empirical footprint performance, line: ideal performance.

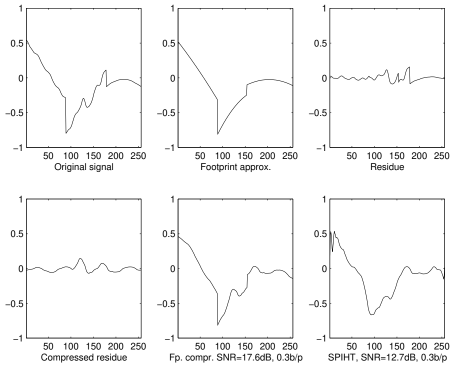

Figure 4: Compression of a piecewise smooth function.

the first k coefficients of the wavelet decomposition are quantized, while the others are set to zero (linear approximation). The allocation of the bits between the piecewise polynomial signal and the residual and the number k of wavelet coefficients that are quantized is chosen off-line, using some a-priori knowledge of the signal. In Figure 4, we show an example of the performance of the proposed compression scheme and compare it with a 1-D version of SPIHT [16]. The signal to compress is given by the union of smooth pieces. In this example, our system outperforms SPIHT by more than 4dB. Since SPIHT is more suited to compress 2-D signals, this comparison is only indicative. However, it shows that a compression system based on footprints can outperform traditional wavelet methods also in the case of piecewise smooth signals.

4.3 Denoising of images with directional wavelets

In images, the standard denoising process is usually done by thresholding the coefficients obtained by the wavelet transform along horizontal and vertical directions. For this reason traditional denoising algorithms do not manage to catch most of the two-dimensional interdependencies present in images. However, such an approach has the main advantage of simplicity. In our method we retain this simplicity, but we manage to better exploit the two-dimensional characteristics of an image. Our algorithm is in spirit similar to cycle-spinning.

We consider several different directions. For each direction, we apply the corresponding directional wavelet transform and hard-threshold the wavelet coefficients. We then invert the wavelet transform to obtain a denoised version of the image. Finally, we average all this different denoised versions to obtain the final image.

Figure 5 shows an example of denoising of the image 'Cameraman'. The original image is affected by additive Gaussian noise and both methods are applied. We can see that the new method which uses, in this case, twenty different directions largely outperforms the standard method.

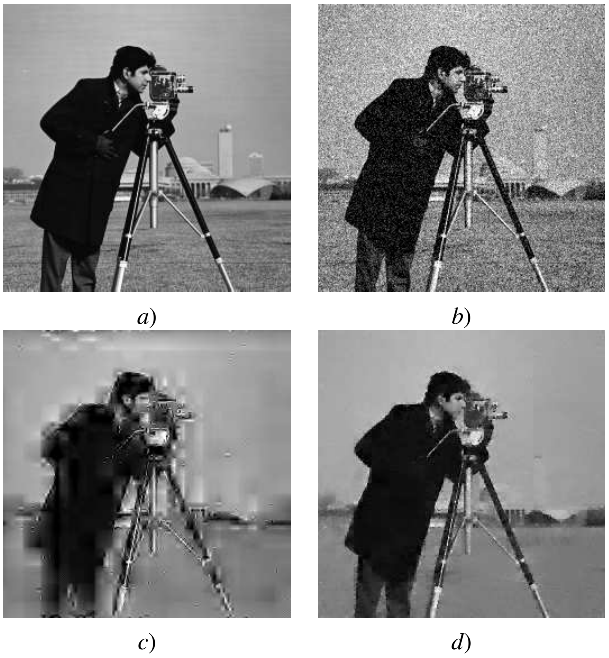

Figure 5: The image 'Cameraman'. (*a*) Original image, (*b*) noised version ($PSNR = 19dB$), (*c*) denoised by the standard method ($PSNR$ = 20.54dB), (*d*) denoised by the new method ($PSNR$ = 25.04dB).

5 Conclusions

In this paper, we have presented new ways to represent signals and images. First, we have introduced the notion of footprints and have shown that footprints provide sparse representations of piecewise smooth signals. This is useful in several signal processing tasks and numerical simulations confirm that footprints outperform wavelet methods in several applications. We have then proposed a simple yet effective way to represent 2-D signals, that is, images. Our new method calculates separable wavelet transform along sets of different directions. Such an approach takes into account two-dimensional interdependencies among discontinuities in images better than the standard method. Application of the new method is possible in various areas of image processing and good results have been shown in denoising.

References

[1] D.L. Donoho, M. Vetterli, R.A. DeVore, and I. Daubechies, "Data compression and harmonic analysis," *IEEE Trans. on Information Theory*, vol. 44(6), pp. 2435–2476, October 1998.

[2] E. J. Candès and D. L. Donoho, "Curvelets – a suprisingly effective nonadaptive representation for objects with edges," in *Curve and Surface Fitting*, A. Cohen, C. Rabut, and L. L. Schumaker, Eds., Saint-Malo, 1999, Vanderbilt University Press.

[3] R.H. Bamberger and M.J.T. Smith, "A filter bank for the directional decomposition of images: Theory and design," *IEEE Trans. Signal Processing*, pp. 882–893, April 1992.

[4] Minh Do, *Directional Multiresolution Image Representations*, PhD, dissertation, Swiss Federal Institute of Technology, November 2001.

[5] R.A. Zuidwijk, "Directional and time-scale wavelet analysis," *SIAM Journal on Mathematical Analysis*, vol. 31(2), pp. 416–430, 2000.

[6] S. Mallat, *A Wavelet Tour of Signal Processing*, Academic Press, 1998.

[7] P.L. Dragotti, *Wavelet footprints and frames for signal processing and communication*, PhD, dissertation, Swiss Federal Institute of Technology, April, 2002.

[8] D.L. Donoho and I.M. Johnstone, "Ideal spatial adaptation via wavelet shrinkage,," *Biometrika*, vol. 81, pp. 425–455, December 1994.

[9] R. R. Coifman and D.L. Donoho, "Translation invariant denoising," *Technical Report 475, Dept. of Statistics, Stanford University*, May 1995.

[10] M. Vetterli, "Wavelets approximation and compression," *IEEE Signal Processing Magazine*, vol. 18, pp. 59–73, September 2001.

[11] A. Cohen, I. Daubechies, O. Guleryuz, and Orchard M.T., "On the importance of combining wavelet-based non-linear approximation with coding strategies," *IEEE Trans. on Information Theory, to appear*.

[12] J.M. Shapiro, "Embedded image coding using zerotrees of wavelets coefficients," *IEEE Trans. on Signal Processing*, vol. 41, pp. 3445–3462, December 1993.

[13] P. Prandoni and M. Vetterli, "Approximation and compression of piecewise smooth functions," *Phil. Trans. Royal Society London*, August 1999.

[14] M. Crouse, R.D. Nowak, and R.G. Baraniuk, "Wavelet-based signal processing using hidden Markov models," *IEEE Trans. Signal Processing*, vol. 2, pp. 886–902, April 1998.

[15] J.E. Bresenham, "Algorithms for computer control of a digital plotter," *IBM Systems Journals*, vol. 4(1), pp. 25–30, 1965.

[16] A. Said and W.A. Pearlman, "A new, fast, and efficient image codec based on set partitioning in hierarchical trees," *IEEE Trans. Circuits and Systems for Video Technology*, vol. 6, no. 3, pp. 243–249, June 1996.