LECTURES ON WAVELET PACKET ALGORITHMS

Mladen Victor Wickerhauser

Department of Mathematics
Washington University
St. Louis, Missouri 63130

April 22, 1992

0. Outline

We begin by defining continuous wavelet packets on **R**. These are square-integrable functions with prescribed smoothness and other properties, which we shall develop to establish the main notions. Our construction will be directed toward numerical applications, so we will restrict ourselves to the quadrature mirror filter algorithm.

Next we will define several discrete algorithms and explore their advantages and disadvantages. We will show the correspondence between wavelet packets and coefficients computed from sampled signals, and relate the convergence of this approximation to the smoothness of the signal. We will define information cost functions and the "best-basis" method. We will count operations and consider practical matters like the memory requirements of the algorithms, periodizing, the spreading of the support of aperiodic wavelet packets, and the combinatorics of constructing wavelet packet bases of increasing generality.

In parallel, we will develop smooth orthogonal local trigonometric transforms. These are properly considered transposes of wavelet packet methods, or alternatively conjugates of wavelet packet methods by the Fourier transform. We will describe both continuous and

Supported in part by the Air Force Office of Scientific Research and the Defense Advanced Research Projects Agency. Copyright 1991 M. V. Wickerhauser; all rights reserved. .

Typeset by $\mathcal{A}_{\mathcal{M}}\mathcal{S}$-TEX

discrete local cosine transforms, and an adaptive local cosine transform useful for signal segmentation.

We will examine several compression methods, both linear and nonlinear. Linear methods include uniform and nonuniform quantization. Nonlinear methods include discarding small coefficients, coalescing to the center of energy within bands, and Karhunen–Loéve methods. We will examine the peculiarities of each method, and discuss the errors in the lossy versions of these algorithms. This will illustrate the relative advantages of wavelet packets, windowed Fourier transforms, and wavelet bases.

We will then generalize to multidimensions by separation of variables. We will explore the combinatorics of higher dimensional wavelets. We will identify matrices with two-dimensional signals or "pictures," and we will show how each picture compression algorithm yields a nonstandard matrix multiplication algorithm.

As a demonstration of the analytic power of best-basis methods, we will perform an automatic analysis of a few canonical signals in the phase plane. The signals will be decomposed into as precise a set of modulated lumps as the Heisenberg uncertainty principle allows, and the product of the analysis will be displayed in an intuitively satisfying manner.

Finally, we will produce several fast numerical algorithms driven by the best-basis method. Among these will be fast approximate Karhunen–Loéve factor analysis, signal segmentation in time and frequency, feature-preserving encryption, and matrix multiplication.

1. Definition of Wavelet Packets

Wave packets. Roughly speaking, a *wave packet* ψ is a square integrable modulated waveform, well localized in both position and frequency. A musical note is an example. It may be assigned three parameters: frequency, scale, and position. The first and third may be taken to be the centers of mass of $|\psi|^2$ and $|\hat{\psi}|^2$, where $\hat{\psi}$ is the Fourier transform of ψ. The second might be taken to be a characteristic width of $|\psi|^2$, or equivalently the uncertainty in the position. By Heisenberg's principle, it is also the reciprocal of the uncertainty in frequency.

Examples of modulated waveforms are easy to construct. Let ϕ be any "sufficiently nice" function and define the modulation, dilation, and translation operators by $\mu_f\phi(t) = e^{ift}\phi(t)$, $\delta_s\phi(t) = s^{1/2}\phi(st)$, and $\tau_p\phi(t) = \phi(t-p)$, respectively. Then the collection of dilated, translated, and modulated ϕ's forms a family of wave packets with parameters

f, s, p. These transformations conserve energy, so the waveforms can be normalized to be unit vectors in L^2. The component of a function x at f, s, p is the inner product of x with the modulated waveform whose parameters are f, s, p. If it is large, we may conclude that x has considerable energy at scale s near frequency f and position p.

The Balian–Low phenomenon [DJJ] prevents the modulated bumps above from being an orthogonal basis. We circumvent this obstruction in two ways. We may either replace the exponentials e^{ift} by one of $\sin \pi(f + \frac{1}{2})t$ or $\cos \pi(f + \frac{1}{2})t$, or we may replace the single ϕ with a family of bumps. The first solution was found by Daubechies, Jaffard, and Journé [DJJ], and later generalized by Laeng [L], Aucher [Au], and Coifman and Meyer [CM2]. A similar solution was found independently by Malvar [Ml]. We will discuss the latest version of this local trigonometric basis in a later chapter. The second method is equivalent to the following construction, which was first noted by Strömberg [S], developed by Mallat [Ma], Daubechies, Grossman and Meyer [DGM] and [D], then generalized into its present form by Coifman, Meyer and Wickerhauser in [CM1] and [CW].

We introduce a new class of orthonormal bases of $L^2(\mathbf{R}^n)$ by constructing a "library" of modulated waveforms out of which various bases can be extracted. In particular, the wavelet basis, the Walsh functions, and rapidly oscillating "wavelet packet" bases are obtained. It is the orthogonality properties that distinguish these new functions from wave packets.

We shall use the notation and terminology of [D], and assume the results therein. Let $h = \{h_n\}$ be an exact quadrature mirror filter satisfying the conditions of Theorem (3.6) in [D], p. 964, i.e.,

$$\sum h_{n-2k} h_{n-2\ell} = \delta_{k\ell}; \qquad \sum h_n = \sqrt{2}.$$

Let $g_k = (-1)^k h_{1-k}$ and define the operations F_0 and F_1 from $\ell^2(\mathbf{Z})$ into "$\ell^2(2\mathbf{Z})$"

$$\begin{aligned} F_0\{s_k\}(j) &= \sum_k s_k h_{k-2j} \\ F_1\{s_k\}(j) &= \sum_k s_k g_{k-2j} \end{aligned} \tag{1.0}$$

These operations have adjoints F_i^* defined by

$$\begin{aligned} F_0^*\{s_j\}(k) &= \sum_j s_j h_{k-2j} \\ F_1^*\{s_j\}(k) &= \sum_j s_j g_{k-2j} \end{aligned} \tag{1.0*}$$

The maps $F_i F_i^*$ are orthogonal projections on $\ell^2(\mathbf{Z})$. The map $\mathbf{F}(s_k) = F_0(s_k) \oplus F_1(s_k) \in \ell^2(\mathbf{Z}) \oplus \ell^2(\mathbf{Z})$ is orthogonal and

$$F_0^* F_0 + F_1^* F_1 = I \tag{1.1}$$

Fixed-scale wavelet packets. We now define a sequence of functions recursively.

$$\begin{aligned} W_{2n}(t) &= \sqrt{2}\sum h_k W_n(2t-k), \\ W_{2n+1}(t) &= \sqrt{2}\sum g_k W_n(2t-k). \end{aligned} \tag{1.2}$$

The function $W_0(t)$ can be identified with the function φ in [D], and W_1 can be identified with the function ψ.

Let us define $m_0(\xi) = \frac{1}{\sqrt{2}}\sum h_k e^{-ik\xi}$ and

$$m_1(\xi) = -e^{i\xi}\bar{m}_0(\xi+\pi) = \frac{1}{\sqrt{2}}\sum g_k e^{ik\xi}$$

Remark. The quadrature mirror condition on the operation $\mathbf{F} = (F_0, F_1)$ is equivalent to the following matrix being unitary:

$$\mathcal{M} = \begin{bmatrix} m_0(\xi) & m_1(\xi) \\ m_0(\xi+\pi) & m_1(\xi+\pi) \end{bmatrix} \tag{1.3}$$

The Fourier transform of Eq. (1.2) when $n = 0$ gives the relations

$$\hat{W}_0(\xi) = m_0(\xi/2)\hat{W}_0(\xi/2)$$

i.e.,

$$\hat{W}_0(\xi) = \prod_{j=1}^{\infty} m_0(\xi/2^j)$$

and

$$\hat{W}_1(\xi) = m_1(\xi/2)\hat{W}_0(\xi/2) = m_1(\xi/2)m_0(\xi/4)m_0(\xi/2^3)\cdots$$

More generally, the relations Eq.(1.2) are equivalent to

$$\hat{W}_n(\xi) = \prod_{j=1}^{\infty} m_{\varepsilon_j}(\xi/2^j) \tag{1.4}$$

and $n = \sum_{j=1}^{\infty} \varepsilon_j 2^{j-1}$, where $\varepsilon_j = 0$ or 1.

We can rewrite Eq. (1.2) as follows.

$$\begin{aligned} W_{2n}(t-\ell) &= \sqrt{2}\sum h_{j-2\ell} W_n(2t-j) = F_0\{\sqrt{2}W_n(2t-j)\}(\ell) \\ W_{2n+1}(t-\ell) &= \sqrt{2}\sum g_{j-2\ell} W_n(2t-j) = F_1\{\sqrt{2}W_n(2t-j)\}(\ell) \end{aligned} \tag{1.5}$$

where $\sqrt{2}W_n(2t-j)$ is viewed as a sequence in j for (t,n) fixed. Using Eq. (1.1) we find:

$$W_n(t-j) = \frac{1}{\sqrt{2}} \sum_i h_{j-2i} W_{2n}\left(\frac{t}{2} - i\right) + \frac{1}{\sqrt{2}} \sum_i g_{j-2i} W_{2n+1}\left(\frac{t}{2} - i\right) \tag{1.6}$$

In the case $n = 0$ we obtain:

$$W_0(t-k) = \frac{1}{\sqrt{2}} \sum_i h_{k-2i} W_0\left(\frac{t}{2} - i\right) + \frac{1}{\sqrt{2}} \sum_i g_{k-2i} W_1\left(\frac{t}{2} - i\right) \tag{1.7}$$

from which we deduce the usual decomposition of a function f in the space Ω_0 (V_0 in [D]), i.e., a pair of series for the function f of the form

$$\begin{aligned} f(t) &= \sum_k s_k^0 W_0(t-k) \\ &= \frac{1}{\sqrt{2}} \sum_i (\sum_k s_k^0 h_{k-2i}) W_0\left(\frac{t}{2} - i\right) + \frac{1}{\sqrt{2}} \sum_i (\sum_k s_k^0 g_{k-2i}) W_1\left(\frac{t}{2} - i\right) \\ &= \sum_i F_0(s_k^0)(i) \frac{1}{\sqrt{2}} W_0\left(\frac{t}{2} - i\right) + \sum_i F_1(s_k^0)(i) \frac{1}{\sqrt{2}} W_1\left(\frac{t}{2} - i\right) \end{aligned}$$

More generally, if we define the linear span of integer translates of W_n's

$$\Omega_n = \{f : f = \sum \omega_k W_n(t-k)\}, \tag{1.8}$$

then we can put

$$f(t) = \sum_i F_0(\omega_k)(i) \frac{1}{\sqrt{2}} W_{2n}\left(\frac{t}{2} - i\right) + \sum_i F_1(\omega_k)(i) \frac{1}{\sqrt{2}} W_{2n+1}\left(\frac{t}{2} - i\right) \tag{1.9}$$

or

$$\sqrt{2} f(2t) = h + g \qquad \text{for } h \in \Omega_{2n} \text{ and } g \in \Omega_{2n+1}.$$

Notice that $\{W_n(t-j) : j \in \mathbf{Z}\}$ is an orthonormal basis for the space Ω_n

We now prove

Theorem 1.1. *The functions $W_n(t-k)$, for integers k, n with $n \geq 0$, form an orthonormal basis of $L^2(\mathbf{R})$.*

Proof. We proceed by induction on n, assuming that the set of translates $\{W_n(t-k) : k \in \mathbf{Z}\}$ forms an orthonormal set of functions and proving that $\{W_{2n}(t-k) : k \in \mathbf{Z}\} \cup \{W_{2n+1}(t-k) : k \in \mathbf{Z}\}$ forms an orthonormal set.

By assumption $\|\sqrt{2}f(2t)\|_2^2 = \|f(t)\|_2^2 = \sum \omega_k^2$ if $f \in \Omega_n$. From the quadrature mirror condition on (F_0, F_1) we get

$$\sum \omega_k^2 = \sum F_0(\omega_k)(i)^2 + F_1(\omega_k)(i)^2.$$

Since $F_0(\omega_k)(i) = \mu_i$ and $F_0(\omega_k)(i) = \nu_i$ can be chosen as two arbitrary sequences of ℓ^2 (arising from $\omega = F_0^*(\mu) + F_1^*(\nu)$) it follows that

$$\int |\sum_i \mu_i W_{2n}(t-i) + \sum_i \nu_i W_{2n+1}(t-i)|^2 = \sum_i \mu_i^2 + \sum_i \nu_i^2$$

which is equivalent to $\{W_{2n}(t-j) : j \in \mathbf{Z}\} \cup \{W_{2n+1}(t-j) : j \in \mathbf{Z}\}$ being an orthonormal set of functions.

Having proved the inductive step, we now observe after k applications that $\{W_0(t-j) : j \in \mathbf{Z}\} \cup \cdots \cup \{W_{2^k-1}(t-j) : j \in \mathbf{Z}\}$ is an orthonormal set of functions. In other words, the subspaces $\Omega_0 \oplus \cdots \oplus \Omega_{2^k-1}$ are mutually orthogonal in $L^2(\mathbf{R})$.

Let us now define $\delta f(t) = \sqrt{2}f(2t)$. Equation (1.8) shows that $\delta\Omega_n = \Omega_{2n} \oplus \Omega_{2n+1}$, or

$$\begin{aligned}
\delta\Omega_0 - \Omega_0 &= \Omega_1, \\
\delta^2\Omega_0 - \delta\Omega_0 = \delta\Omega_1 &= \Omega_2 \oplus \Omega_3, \\
\delta^3\Omega_0 - \delta^2\Omega_0 = \delta\Omega_2 \oplus \delta\Omega_3 &= \Omega_4 \oplus \Omega_5 \oplus \Omega_6 \oplus \Omega_7,
\end{aligned}$$

and in general,

$$\delta^k\Omega_0 - \delta^{k-1}\Omega_0 = \Omega_{2^{k-1}} \oplus \Omega_{2^{k-1}+1} \cdots \oplus \Omega_{2^k-1}. \tag{1.10}$$

These may be telescoped to give the decomposition

$$\delta^k\Omega_0 = \Omega_0 \oplus \Omega_1 \oplus \cdots \oplus \Omega_{2^k-1}. \tag{1.11}$$

Since $\delta^k\Omega_0 \to L^2(\mathbf{R})$ as $k \to \infty$, (see [D]), it follows that $\{W_n(t-j) : n \geq 0, j \in \mathbf{Z}\}$ is complete. □

The collection $\{W_n(t-j)\}$ is a wavelet packet analog of Walsh functions. If we use the Haar filters $h_1 = h_0 = 1/\sqrt{2}$ and restrict our functions to $0 \leq t < 1$, we note that $\{W_n(t)\}$ is the Walsh functions in Paley or "natural" order.

General wavelet packets. All of the functions W_n in the preceding section have a fixed scale, but we observe that mixed-scale decompositions of L^2 are also possible. From Eq. (1.9) we observe that $\delta\Omega_n = \Omega_{2n} \oplus \Omega_{2n+1}$, or more generally that

$$\delta^k\Omega_n = \Omega_{2^k n} \oplus \cdots \oplus \Omega_{2^k(n+1)-1}. \tag{1.12}$$

This allows us to refine the decomposition $L^2 = \oplus_n \Omega_n$ by scales as embodied in the following:

Theorem 1.2. *For every partition P of the nonnegative integers into sets of the form $I_{kn} = \{2^k n, \ldots, 2^k(n+1)-1\}$, the collection of functions $\{2^{k/2}W_n(2^k t - j) : I_{kn} \in P, j \in \mathbf{Z}\}$ is an orthonormal basis for $L^2(\mathbf{R})$.*

Proof. We note that $\{2^{k/2}W_n(2^k t - j) : j \in \mathbf{Z}\}$ is an orthonormal basis for $\delta^k \Omega_n$. If $\cup_P I_{kn}$ is all of the nonnegative integers, then Eq. (1.12) implies $\sum_{I_{nk} \in P} \delta^k \Omega_n = \oplus_j \Omega_j = L^2(\mathbf{R})$. □

Remark. We may also think of I_{nk} as the dyadic subinterval $[2^{-k}n, 2^{-k}(n+1)[$ of $[0,1[$. Such an indexing convention gives a faithful correspondence between disjoint dyadic decompositions of $[0,1[$ and orthonormal wavelet packet subsets of L^2.

Definition. *A* wavelet packet basis *of $L^2(\mathbf{R})$ is any orthonormal basis selected from among the functions $2^{k/2}W_n(2^k t - j)$.*

Beside the Walsh-type basis $\Omega_0 \oplus \Omega_1 \oplus \cdots \oplus \Omega_k \oplus \ldots$, examples of wavelet packet bases include the *wavelet basis* $\Omega_0 \oplus \Omega_1 \oplus \delta\Omega_1 \oplus \cdots \oplus \delta^k \Omega_1 \oplus \ldots$ and the *subband basis* $\delta^K \Omega_0 \oplus \delta^K \Omega_1 \oplus \cdots \oplus \delta^K \Omega_n \oplus \ldots$.

Influence of the QMFs. The choice of quadrature mirror filter influences the properties of wavelet packets, including smoothness and number of vanishing moments. The computational complexity of the wavelet packet algorithm grows linearly with filter length.

Frequency resolution of wavelet packets.. Suppose we set

$$m_0(\xi) = \sum_{k=-\infty}^{\infty} 1_{[-\frac{\pi}{2}, \frac{\pi}{2}]}(\xi - 2\pi k)$$

which is periodic on $\mathbf{R}$ with period 2π. Likewise, set $m_1(\xi) = 1 - m_0(\xi)$. The pair m_0, m_1 satisfies the orthogonality condition (1.3), and corresponds to infinitely long QMF's. We can explicitly compute the limits $\hat{W}_0(\xi) = 1_{[-\pi,\pi]}(\xi)$ and $\hat{W}_1(\xi) = 1_{[-2\pi,-\pi]}(\xi) + 1_{[\pi,2\pi]}(\xi)$.

The Gray code permutation $f \mapsto f'$ is $f_i = f'_{i+1} + f'_i \pmod 2$, where f_i is the ith binary digit of f. It relates Paley order to sequency order for Walsh functions, and appears naturally in the frequency localization of wavelet packets.

Theorem 1.3. *For the special m_0, m_1 above, we have $\hat{W}_n(\xi) = 1_{[j\pi,(j+1)\pi]}(|\xi|)$, where n is the Gray-code permutation of j.*

Proof. Since this is true for $n = j = 0$ and $n = j = 1$, let us prove the inductive step.

Now $m_0(\xi/2) = \sum_k 1_{[-\frac{\pi}{2}, \frac{\pi}{2}]}(\frac{\xi}{2} - 2\pi k) = \sum_k 1_{[-\pi,\pi]}(\xi - 4\pi k) = \sum_k 1_{[(4k-1)\pi,(4k+1)\pi]}(\xi)$, which is periodic with period 4π. Also, $\hat{W}_n(\xi/2) = 1_{[2j\pi,2(j+1)\pi]}(|\xi|)$. Therefore,

$$\hat{W}_{2n}(\xi) = m_0(\xi/2)\hat{W}_n(\xi/2) = \begin{cases} 1_{[2j\pi,(2j+1)\pi]}(|\xi|), & \text{if } j \text{ is even,} \\ 1_{[(2j+1)\pi,2(j+1)\pi]}(|\xi|), & \text{if } j \text{ is odd.} \end{cases}$$

Likewise, $m_1(\xi/2) = \sum_k 1_{[\frac{\pi}{2},\frac{3\pi}{2}]}(\frac{\xi}{2} - 2\pi k) = \sum_k 1_{[\pi,3\pi]}(\xi - 4\pi k) = \sum_k 1_{[(4k+1)\pi,(4k+3)\pi]}(\xi)$, so that

$$\hat{W}_{2n+1}(\xi) = m_1(\xi/2)\hat{W}_n(\xi/2) = \begin{cases} 1_{[2j\pi,(2j+1)\pi]}(|\xi|), & \text{if } j \text{ is odd,} \\ 1_{[(2j+1)\pi,2(j+1)\pi]}(|\xi|), & \text{if } j \text{ is even.} \end{cases}$$

Thus the wavelet packet whose Fourier transform is $1_{[2j\pi,(2j+1)\pi}(|\xi|)$ will be indexed by $2n$ if j is even and by $2n+1$ if j is odd. If n is the Gray-code of j, then this index is the Gray-code of $2j$. The result for $2j+1$ is left as an exercise. □

The above result generalizes to finite impulse response filters by a somewhat more technical argument, which may be seen in [CMW]. We avoid these technicalities so as not to obscure the main point. An interesting difference in all the compactly supported cases except Haar is that the ℓ^∞ norm of the nth wavelet packet is not bounded. It grows no faster than n^δ, for some $0 < \delta \leq 1/4$. This implies that the support of the Fourier transform of smooth wavelet packets spreads like n^δ as n increases, making the frequency localization less and less precise. The practical consequences of such spreading can be severe, but are easily avoided by controlling n. We prevent the number of sample points within a window from exceeding a fixed limit determined by the required frequency precision.

2. Discrete Wavelet Packets

Discrete measurements of functions in $L^2(\mathbf{R})$ must be regarded as averages. If the function is continuous, then averages over small intervals can be approximated by evaluations at points inside the intervals. The disagreement between these two pictures is controlled by the regularity of the signal and the averaging function.

Approximating functions by sampling. Suppose that $x = x(t)$ is a square integrable function. If it is band-limited—i.e., if there is a finite $K > 0$ such that $\hat{x}(\xi) = 0$ if $|\xi| > K$—then Shannon's sampling theorem ([SW], p.53) asserts that

$$x(t) = \sum_{n=-\infty}^{\infty} x(\frac{n}{2K}) \frac{\sin \pi(2Kt - n)}{\pi(2Kt - n)}$$

Namely, x is determined exactly at all points by its values at discrete points spaced $1/2K$ apart. These values are well defined because band-limited functions are smooth. All signals are band-limited functions in practice.

The function $\sin \pi(2Kt - n)/\pi(2Kt - n)$ may be called a Shannon wavelet packet. It corresponds to a function perfectly band limited to the interval $[-K, K]$. Although it is not compactly supported—in fact, it has quite poor decay—it is useful for estimating the frequency localization of better behaved wavelet packets.

If the signal x is periodic and band-limited to $(-K, K)$, then x is determined exactly by at most $2K$ samples in a period. This is the trivial observation that a polynomial of degree $2K - 1$ is determined by its values at any $2K$ distinct points.

Computing averages from samples. Averages of a band-limited signal may be computed exactly from samples. Let $\phi = \phi(t)$ be any smooth integrable (hence square-integrable) averaging function; this means that ϕ has unit mass. Then

$$\langle x, \phi \rangle = \int x(t)\phi(t)\, dt = \sum_{n=-\infty}^{\infty} x(n/2K) \int \phi(t) \frac{\sin \pi(2Kt - n)}{\pi(2Kt - n)}\, dt = \sum_{n=-\infty}^{\infty} a_n x(n/2K),$$

where $a_n = \int \phi(t) \frac{\sin \pi(2Kt-n)}{\pi(2Kt-n)}\, dt$. The sequence $\{a_n\}$ is square-summable, because the functions $\frac{\sin \pi(2Kt-n)}{\pi(2Kt-n)}$ are orthonormal in $L^2(\mathbf{R})$. We can evaluate a_n by Plancherel's formula:

$$a_n = \int \phi(t) \frac{\sin \pi(2Kt - n)}{\pi(2Kt - n)}\, dt = \frac{1}{2K} \int_{-K}^{K} e^{2\pi i \frac{n}{2K} \xi} \hat{\phi}(\xi)\, d\xi.$$

This gives the following:

Proposition 2.1. *If the functions x and ϕ are band-limited to $(-K, K)$, then*

$$\langle x, \phi \rangle = \frac{1}{2K} \sum_{n=-\infty}^{\infty} x(n/2K)\phi(n/2K).$$

Proof. Since $\hat{\phi}(\xi) = 0$ if $|\xi| > K$, we can evaluate

$$a_n = \frac{1}{2K} \int_{-K}^{K} e^{2\pi i \frac{n}{2K} \xi} \hat{\phi}(\xi)\, d\xi = \frac{1}{2K} \int_{-\infty}^{\infty} e^{2\pi i \frac{n}{2K} \xi} \hat{\phi}(\xi)\, d\xi = \frac{1}{2K} \phi(n/2K) \quad \square$$

Wavelet scaling functions. Let $x = x(t)$ be a function in $L^2(\mathbf{R})$, and let $\{x_p : p \in \mathbf{Z}\}$ be the coefficients of its projection onto $\delta^L \Omega_0$. Then $x_p = \int 2^{L/2} W_0(2^L t - p) x(t)\, dt$. The L^2 function given by this projection will be denoted $P_L x(t) = \sum_p x_p 2^{L/2} W_0(2^L t - p)$. The error in this approximation is controlled by the projection of x onto the orthogonal complement $\delta^L \Omega_0^\perp$ in L^2, which we may denote by the function $(I - P_L)x$.

This projection is small if both x and W_0 are regular. We have the following results from Littlewood–Paley theory:

Proposition 2.2. *If both x and W_0 have d continuous derivatives, then $\|(I - P_L)x\|_2 < C2^{-dL}\|x\|_2$.* □

Proposition 2.3. *If both x and W_0 are of class $C^d(\mathbf{R})$, and $\int t^k W_0(t)\,dt = 0$ for all $0 < k < d$, then $\|x - 2^{L/2}x_p\|_* < C\|x\|_* 2^{-dL}$, where $\|x\|_* = \sup_{p<2^L t<p+1} |x(t)|$.* □

Proofs of these facts may be found in the excellent two-volume monograph by Meyer [Me].

Numerical calculation of wavelet packet coefficients. Now let $\{x_p\}$ be the coefficients of $x(t)$ in $\delta^L\Omega_0$. From these we may calculate the coefficients of $x(t)$ in any space $\delta^k\Omega_n$, for $0 \le k \le L$ and $0 \le n < 2^{L-k}$, by applying the functions F_0 and F_1 to the sequence $\{x_p\}$. We may introduce the notation

$$x_p^{fs} = \int_{-\infty}^{\infty} 2^{s/2} W_f(2^s t - p)x(t)\,dt \qquad p \in \mathbf{Z}; 0 \le s \le L; 0 \le f < 2^{L-s} \tag{2.1}$$

The coefficients of $x(t)$ in the subspace $\delta^k\Omega_n$ form the sequence $\{x_p^{nk} : p \in \mathbf{Z}\}$. We obtain the following:

Lemma 2.4. *The coefficient sequences $\{x_p^{fs}\}$ satisfy the recursion relations*

$$x_p^{2f,s-1} = F_0\{x^{fs}\}(p)$$
$$x_p^{2f+1,s-1} = F_1\{x^{fs}\}(p)$$

Proof. Transpose the defining formulas for the wavelet packets W_n onto the coefficient sequences. □

Lemma 2.5. *If $0 \le s \le L$ and $0 \le f < 2^{L-s}$, then $x_n^{fs} = F_{\varepsilon_1} \dots F_{\varepsilon_{L-s}}\{x_p\}(n)$, where ε_i is the ith binary digit of f.*

Proof. This is a straightforward induction on Lemma 2.4. □

Lemma 2.5 gives a fast construction for all of the inner products x_p^{fs}. Each application of F_i costs $O(N)$ multiply-adds, where N is the number of signal samples.

Libraries of rapidly constructible functions. Wavelet packets form a library of functions. There are infinitely many of them in the continuum limit, but only finitely many in the approximation $\delta^L\Omega_0$, restricted to a compact interval. These we can approximate by applying the adjoint of the algorithm to elementary sequences.

Lemma 2.6. *The coefficients of $2^{s/2}W_f(2^st-p)$ in the space $\delta^L\Omega_0$ with respect to its standard basis $\{2^{L/2}W(2^Lt-j) : j \in \mathbf{Z}\}$ are given by $x_j = F^*_{\varepsilon_{L-s}} \dots F^*_{\varepsilon_1}\{1_p\}(j)$, where $f = \sum_{i=1}^{L-s} \varepsilon_i 2^{i-1}$, and 1_p denotes the elementary sequence which has a 1 in the p position and 0's elsewhere.*

Proof. Suppose that the coefficients x_j of x in $\delta^L\Omega_0$ are as claimed. By Lemma 2.5, the projection of x onto $\delta^s\Omega_f$ has coefficients $F_{\varepsilon_1} \dots F_{\varepsilon_{L-s}}\{x_j\} = 1_p$ by the quadrature mirror filter equations $F_iF_i^* = I$. The projection onto $\delta^s\Omega_{f'}$ vanishes, since $F_iF_{i'}^* = 0$ when $i \neq i'$. Since $\delta^s\Omega_0 \oplus \delta^s\Omega_1 \oplus \cdots \oplus \delta^s\Omega_{2^{L-s}-1}$ is a basis for $\delta^L\Omega_0$, we conclude that $x(t) = 2^{s/2}W_f(2^st-p)$. □

The collection $\{2^{s/2}W_f(2^st-p)\}$ forms a library of rapidly constructible functions, with a natural organization deriving from its construction via Lemma 2.6. The paths of filter convolution-decimations F_0 and F_1 form a binary tree, with root $\delta^L\Omega_0$ and leaves $\Omega_0, \dots, \Omega_{2^L-1}$. There is an easy generalization to arbitrary trees with similar orthogonality properties, obtained by using other filters [CW].

A useful picture of the tree of wavelet packet coefficients is that of a rectangle of dyadic blocks. The row number within the rectangle indexes the scale of the wavelet packets listed therein. The column number indexes both the frequency and position parameters. We may choose to group the wavelet packets either by position or by frequency. Grouping by position fills each row of the rectangle with adjacent windowed spectral transforms, analogous to windowed FFT, with the window size determined by the row number and the window position corresponding to the location of the group. The frequency parameter increases within the group.

We will group the coefficients by frequency, since that gives a more efficient implementation, and since the transformation to the other form is obvious. The boxes of coefficients in the rectangle correspond to the decomposition of $\delta^L\Omega_0$ into the subspaces $\delta^k\Omega_n$, for $0 \le k \le L$ and $0 \le n < 2^{L-k}$. The top box corresponds to $\delta^L\Omega_0$, the bottom boxes correspond to Ω_n, for $0 \le n < 2^L$, and box n on level k (counting the bottom as level 0) corresponds to subspace $\delta^k\Omega_n$.

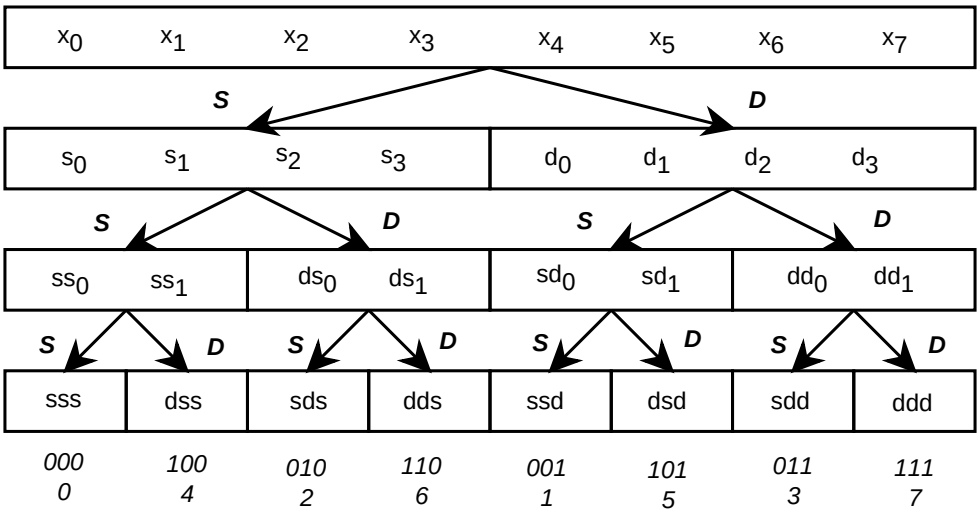
For definiteness, consider a function defined at 8 points $\{x_1, \dots, x_8\}$, i.e., a vector in $\mathbf{R}^8$. Then the (periodized) wavelet packet coefficients of this function look like the following:

A rectangle of wavelet packet coefficients.

x_1	x_2	x_3	x_4	x_5	x_6	x_7	x_8
s_1	s_2	s_3	s_4	d_1	d_2	d_3	d_4
ss_1	ss_2	ds_1	ds_2	sd_1	sd_2	dd_1	dd_2
sss_1	dss_1	sds_1	dds_1	ssd_1	dsd_1	sdd_1	ddd_1

Each row is computed from the row above it by one application of either F_0 or F_1, which we think of as "summing" (s) or "differencing" (d) operations, respectively. Thus, for example, the subblock $\{ss_1, ss_2\}$ comes from the application of F_0 to $\{s_1, s_2, s_3, s_4\}$, while $\{ds_1, ds_2\}$ comes similarly from F_1. The two descendent s and d subblocks on row $n+1$ are determined by their mutual parent on row n, which conversely is determined by them through the adjoint anticonvolutions F_0^* and F_1^*. The process is depicted in this exploded view:

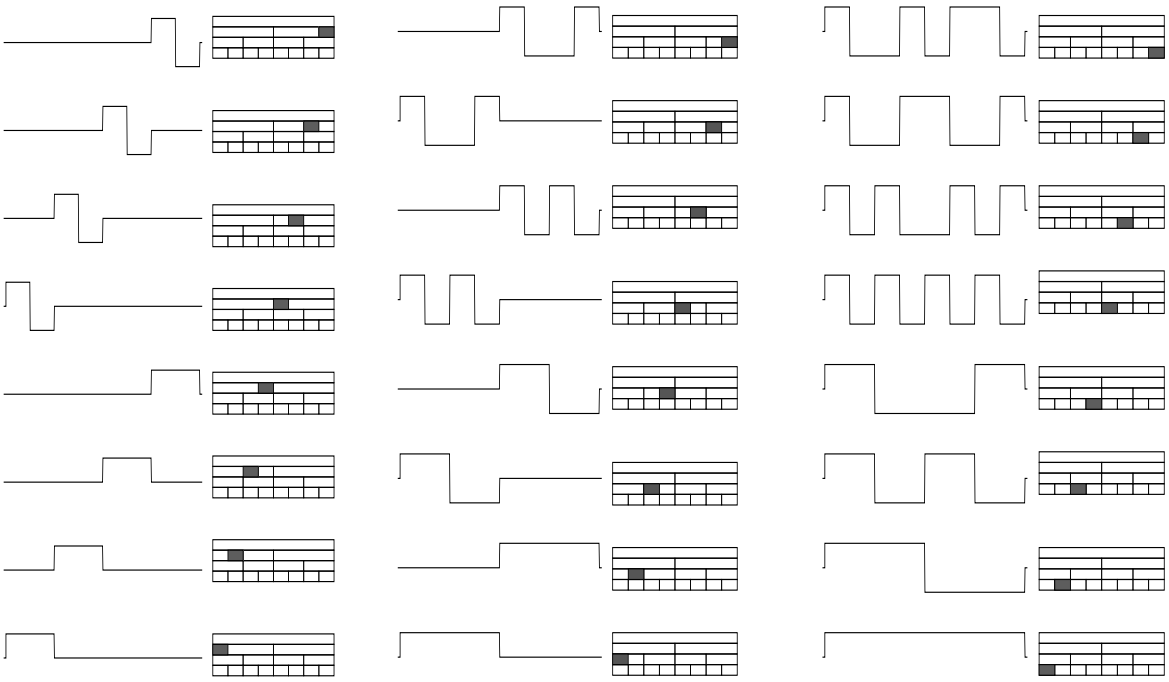
Naturally ordered wavelet packets on $\mathbf{R}^8$



In the simplest case, where we use the Haar filters $h = \{\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}}\}$ and $g = \{\frac{1}{\sqrt{2}}, -\frac{1}{\sqrt{2}}\}$, we have in particular $ss_1 = \frac{1}{\sqrt{2}}(s_1 + s_2)$, $ss_2 = \frac{1}{\sqrt{2}}(s_3 + s_4)$, $ds_1 = \frac{1}{\sqrt{2}}(s_1 - s_2)$, and $ds_2 = \frac{1}{\sqrt{2}}(s_3 - s_4)$. For $N = 8$, we can draw the functions which correspond to the entries in the rectangle. These are displayed in the figures below.

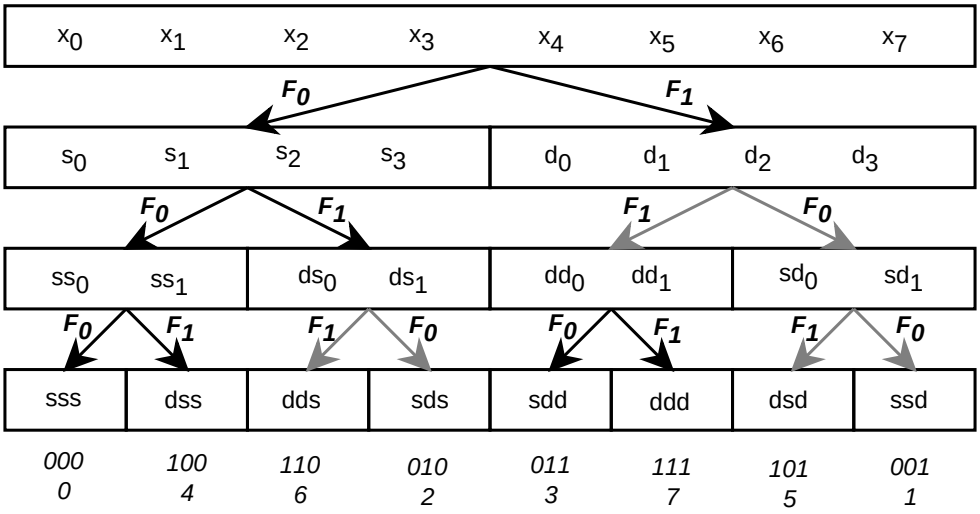
Haar wavelet packets on $\mathbf{R}^8$*:*

Smallest scale = level 1; Intermediate scale = level 2; Largest scale = level 3.



The algorithm produces Haar wavelet packets in the "Paley" or "natural" order. The algorithm may be easily modified to produce "sequency" ordered wavelet packets: what is needed is to exchange F_0 and F_1 whenever the parent's sequency is odd. The diagram below depicts the permuted wavelet packet transform on an 8 point signal, which produces Walsh functions in sequency order if we use the Haar filters above for F_0 and F_1. Sequency has a strict definition only for Walsh functions, where it is the number of zero-crossing of a function which takes only the values 1 and -1. The nth Shannon wavelet packet, in sequency order, is band-limited to the intervals $\pm[n, n+1[$. If we define the appropriate notion of "main frequency" in the intermediate case of smooth, compactly supported wavelet packets, we see that main frequency increases monotonically with sequency order.

Sequency ordered wavelet packets on $\mathbf{R}^8$



Paley order can also be obtained from sequency order by the Gray code permutation.

3. Example Bases of Discrete Wavelet Packets

From the rectangle, we may choose subsets of N coefficients which correspond to orthonormal bases for $\mathbf{R}^N$.

The multiresolution or wavelet decomposition of Mallat [Ma] is a particular descending chain of subspaces Ω in our picture. For example, a discrete wavelet basis $\delta^3\Omega_0 = \Omega_0 \oplus \Omega_1 \oplus \delta^1\Omega_1 \oplus \delta^2\Omega_1$ is the subset corresponding to the labelled boxes in the figure below:

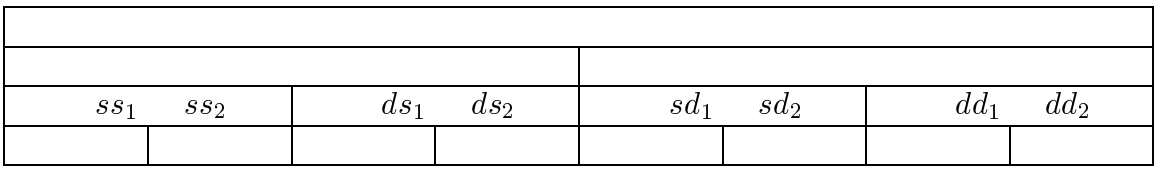
The wavelet basis.

<table>
<tr><td colspan="8"></td></tr>
<tr><td colspan="4"></td><td>d_1</td><td>d_2</td><td>d_3</td><td>d_4</td></tr>
<tr><td colspan="2"></td><td>ds_1</td><td>ds_2</td><td colspan="2"></td><td colspan="2"></td></tr>
<tr><td>sss_1</td><td>dss_1</td><td></td><td></td><td></td><td></td><td></td><td></td></tr>
</table>

Other choices give other orthonormal basis subsets. A single level of the rectangle corresponds roughly to a windowed Fourier transform. In the Haar case described above, the

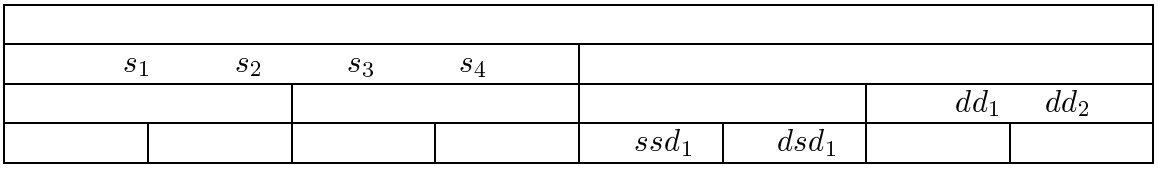
bottom level is exactly the Walsh basis, and intermediate levels are windowed Walsh bases at all dyadic window widths. Longer filters give smoother Walsh-like functions which are analogous to exponentials. The intermediate levels mix the resolvable frequencies together into subbands. For example, the complete level below corresponds to the decomposition $\delta^3\Omega_0 = \delta^1\Omega_0 \oplus \delta^1\Omega_1 \oplus \delta^1\Omega_2 \oplus \delta^1\Omega_3$:

A subband basis.



Any other subset corresponding to a disjoint dyadic will work as well, such as the unnamed basis $\delta^3\Omega_0 = \delta^2\Omega_0 \oplus \Omega_4 \oplus \Omega_5 \oplus \delta^1\Omega_3$

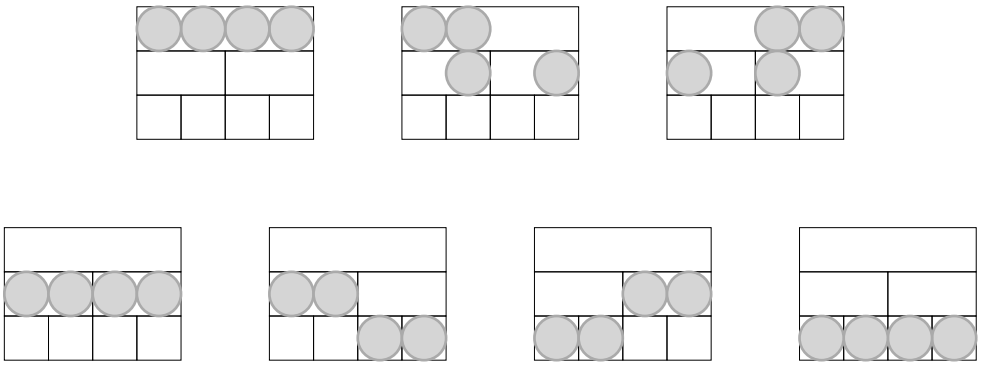
An orthonormal basis subset.



Counting the wavelet packet bases. By Theorem 1.2 and the remark following it, every disjoint dyadic cover of $[0, 1[$ down to widths 2^{-L} corresponds to an orthonormal basis. The number of such covers may be counted by induction. Let A_L be the number of graphs in the coefficient rectangle of $1 + L$ rows. Then $A_0 = 1$ and we have the relation $A_{L+1} = 1 + A_L^2$, which implies that $A_{L+1} > 2^{2^L}$. If our sequences are limited to $N = 2^L$ nonzero coefficients, then this will provide more than 2^N orthogonal bases for R^N.

The bases described by Theorem 1.2 are particularly useful and quite numerous, but they are not all the possible wavelet packet bases. In particular, if we use the Haar filters, then there are wavelet packet bases not corresponding to decompositions into whole blocks, as displayed in the diagram below.

Haar wavelet packet bases for $\mathbf{R}^4$



The two rectangles on the right in the top row do not correspond to disjoint dyadic covers of $[0, 1[$.

4. The Best-Basis Method

We now define a cost function on sequences and search for its minimum over all bases in a library. These functions have practical importance if they measure the cost of describing a sequence. For a given vector, their minima are the most efficient representations, or "best bases."

The super fast growth in L of the number of wavelet packet bases of $\delta^L\Omega_0$ implies that this problem is intractable in general. We will restrict our attention to those functions which split nicely across cartesian products, so that the the search can proceed by divide-and-conquer. This will give a fast algorithm for finding the minimum.

Information cost functions. We can define any sort of real-valued functional $\mathcal{M}$ on sequences, but the most useful are those that measure concentration. By this we mean that $\mathcal{M}$ should be large when the coefficients are roughly the same size and small when all but a few coefficients are negligible. This property should hold on the unit sphere in ℓ^2, since we will be measuring coefficient sequences in various orthogonal bases.

Definition. *A map* $\mathcal{M}$ *from sequences* $\{x_i\}$ *to* $\mathbf{R}$ *is called an* additive information cost function *if* $\mathcal{M}(0) = 0$ *and* $\mathcal{M}(\{x_i\}) = \sum_i \mathcal{M}(x_i)$.

This name is warranted by our goal of describing a vector with the most economical list of coefficients. We will shorten the name to "information cost." Some useful examples of

information costs include:

(1) *Number above a threshold.* Set an arbitrary threshold ϵ and count the elements in the sequence x whose absolute value exceeds ϵ. This is an additive measure of information. It gives the number of coefficients needed to transmit the signal to precision ϵ.

(2) *Concentration in ℓ^p norm, $p < 2$.* Choose an arbitrary $p < 2$ and set $\mathcal{M}(x) = \|\{x\}\|_p$. This is clearly an information cost function, and we observe that the smaller is the ℓ^p norm of a function of energy 1, the more concentrated is its energy into a few coefficients.

(3) *Entropy.* Define the Shannon-Weaver entropy of a sequence $x = \{x_j\}$ by $\mathcal{H}(x) = -\sum_j p_j \log p_j$, where $p_j = \frac{|x_j|^2}{\|x\|^2}$ and we set $p \log p = 0$ if $p = 0$. This is not an additive information cost function. However, the function $\lambda(x) = -\sum_j |x_j|^2 \log |x_j|^2$ is. By the relation $\mathcal{H}(x) = \|x\|^{-2}\lambda(x) + \log \|x\|^2$, minimizing the latter minimizes the former. A classical fact about entropy is that $\exp \mathcal{H}(x)$ is proportional to the number of coefficients needed to represent the signal to a fixed mean square error.

(4) *Logarithm of energy.* Let $\mathcal{M}(x) = \sum_j \log |x_j|^2$, setting $\log 0 = 0$ whenever necessary. This is obviously an additive information cost. It may be interpreted as the entropy of a Gauss-Markov process composed of N uncorrelated Gaussian random variables of variances $\sigma_1^2 = |x_1|^2, \ldots, \sigma_N^2 = |x_N|^2$. The redefinition of log at 0 is equivalent to ignoring any unchanging components in the process. Minimizing this function finds the best approximation to the Karhunen–Loéve basis for the process, which attains the global minimum for $\mathcal{M}$ over the whole orthogonal group.

Theoretical dimension. Suppose that $\{x_n\}$ belongs to both L^2 and $L^2 \log L$. If $x_n = 0$ for all sufficiently large n, then in fact the signal is finite dimensional. Generalizing this notion, we can compare sequences by their rate of decay, i.e., the rate at which their elements become negligible if they are rearranged in decreasing order.

Definition. *The* theoretical dimension *of* $\{x_n\}$ *is*

$$d = \exp(-\sum_n p_n \log p_n)$$

where $p_n = |x_n|^2 \|x\|^{-2}$.

Proposition 4.1. *If $x_n = 0$ for all but finitely many (say N) values of n, then $1 \le d \le N$.*

Proof. In [SW] (as well as in any later information theory text), it is shown that $0 \le -\sum_{n=1}^N p_n \log p_n \le \log N$. □

Proposition 4.2. *If $\{x_n\}$ and $\{x'_n\}$ are rearranged so that both $\{p_n\}$ and $\{p'_n\}$ are monotone decreasing, and if we have $\sum_{0<n<m} p_n \geq \sum_{0<n<m} p'_n$ for all m, then $d \leq d'$.*

Proof. The ordering of partial sums guarantees (by induction) that each p'_n is a convex combination of $0, p_1, \ldots, p_n$. Then convexity of the entropy function yields the result. □

Fast searches for minimum information cost. Additive information cost functions are functionals on the manifold of orthonormal bases, i.e., the orthogonal group. In the finite dimensional case, this manifold is compact, hence there is a global minimum for every continuous information cost. Unfortunately, this minimum will not be a rapidly computable basis in general, nor will the search for a minimum be of low complexity. Therefore, we will restrict our attention to a discrete set of bases sprinkled around the orthogonal group which has the property that each basis has an associated fast transform ($O(N \log N)$ or better) and for which the search for a global minimum of $\mathcal{M}$ converges in $O(N)$ operations, for N-sample signals.

Definition. *A library of orthonormal bases is* a (binary) tree *if it satisfies:*

1. *Subsets of basis vectors can be identified with intervals of* $\mathbf{N}$ *of the form* $I_{nk} = [2^k n, 2^k(n+1)[$, *for* $k, n \geq 0$.
2. *Each basis in the library corresponds to a disjoint cover of* $\mathbf{N}$ *by intervals* I_{nk}.
3. *If* V_{nk} *is the subspace identified with* I_{nk}, *then* $V_{n,k+1} = V_{2n,k} \oplus V_{2n+1,k}$.

Remark. In multidimensions, we must extend this notion to libraries which can be organized as more general trees. This can be done by replacing (3) with the condition that for each $k, n \geq 0$, we have an integer $b > 1$ such that $V_{n,k+1} = V_{bn,k} \oplus \cdots \oplus V_{bn+b-1,k}$. It will not change the subsequent argument.

Let $\mathcal{M}$ be given and let x be a vector in a (separable) space V. Let B denote an orthonormal basis from a library, and denote by Bx the sequence of coefficients of x in basis B.

Definition. *The* best basis *for $x \in V$ relative to $\mathcal{M}$ is that B for which $\mathcal{M}(Bx)$ is minimal.*

If the library is a tree, then we can find the best basis by induction on k. Denote by B_{nk} the basis of vectors corresponding to I_{nk}, and by A_{nk} the best basis for x restricted to the span of B_{nk}. For $k = 0$, there is a single basis available, namely the one corresponding to $I_{n,0}$, which is therefore the best basis: $A_{n,0} = B_{n,0}$ for all $n \geq 0$. We construct $A_{n,k+1}$ for all $n \geq 0$ as follows:

$$A_{n,k+1} = \begin{cases} B_{n,k+1}, & \text{if } \mathcal{M}(B_{n,k+1}x) < \mathcal{M}(A_{2n,k}x) + \mathcal{M}(A_{2n+1,k}x), \\ A_{2n,k} \oplus A_{2n+1,k}, & \text{otherwise.} \end{cases} \tag{4.1}$$

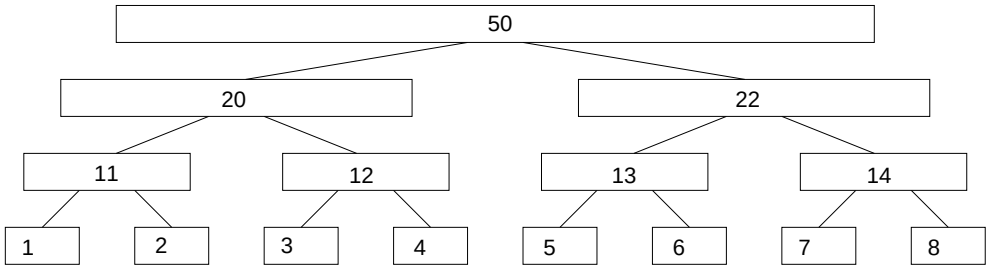
Fix $K \geq 0$ and let V be the span of I_{0K}. We have the following:

Proposition. *The algorithm in Eq.* (4.1) *yields the best basis for x relative to $\mathcal{M}$.*

Proof. This can be shown by induction on K. For $K = 0$, there is only one basis for V. If A' is any basis for $V_{0,K+1}$, then either $A' = B_{0,K+1}$ or $A' = A'_0 \oplus A'_1$ is a direct sum of bases for $V_{0,K}$ and $V_{1,K}$. Let A_0 and A_1 denote the best bases in these subspaces. By the inductive hypothesis, $\mathcal{M}(A_i x) \le \mathcal{M}(A'_i x)$ for $i = 0, 1$, and by Eq. (4.1), $\mathcal{M}(Ax) \le \min\{\mathcal{M}(B_{0,K+1}x), \mathcal{M}(A_0 x) + \mathcal{M}(A_1 x)\} \le \mathcal{M}(A'x)$. □

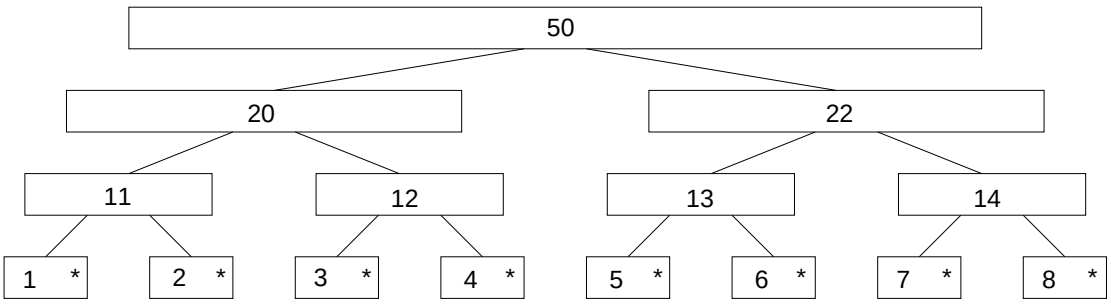
To illustrate the algorithm, consider the following example expansion into a 3-level tree wavelet packets. We have placed numbers representing information costs inside the nodes of the tree:

Sample information costs in a wavelet packet tree



We start by marking all the bottom nodes, as indicated by the asterisks in the diagram. Their total information cost is an initial value which we will try to reduce. This guarantees that the algorithm will terminate:

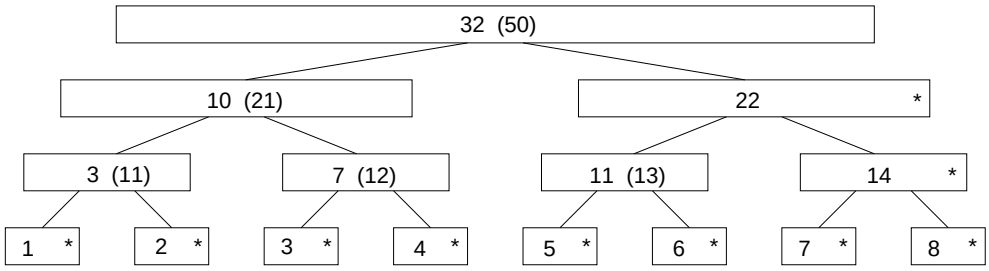
Step 1 in the best basis search: mark all bottom nodes



Whenever a parent node is of lower information cost than the children, we mark the parent. If the children have lower information cost, we do not mark the parent, but we assign the lower total information cost of the children to the parent. This inductive step

prevents our having to examine any node more that twice, once as a child, and once as a parent:

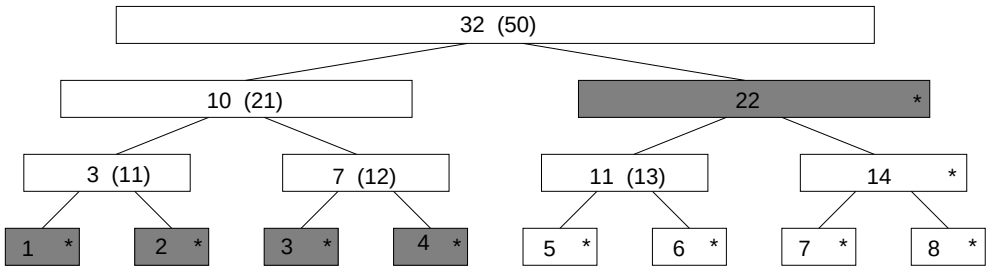
Step 2 in the best basis search: mark all nodes of lower cost



In this diagram, the former information costs of adjusted nodes is displayed in parentheses. This marking may be done at the same time as the coefficient values in the nodes are calculated, if the construction of the tree proceeds in depth-first order. In that case, the information costs of the best basis below each node is known before the recursive descent into that node returns. The algorithm terminates because the tree has finite depth. By induction, the topmost node (or "root") will return the minimum information cost of any basis subset below itself, namely, the best basis.

Finally, after all the nodes have been examined, we take the topmost marked nodes, which constitutes a basis by the preceding theorem. This is equivalent to a depth-first search, and requires no more operations than there are nodes in the tree:

Step 3 in the best basis search: retain topmost marked nodes



The best-basis nodes are shaded in the above diagram. When the topmost marked node is encountered, the remaining nodes in its subtree are pruned away.

Finally, the coefficients in the best-basis nodes are extracted into an output list, in some order. This may of course be done at the same time as the depth first search for topmost

marked nodes. We may assume that they are tagged by their node of origin, i.e., that they come out as quadruples (a, f, s, p) of amplitude, frequency, scale, and position. From the original signal, therefore, to this tagged list takes $O(N \log N)$ operations.

5. Local Trigonometric Bases

Wavelet packets are not the only library useful for signal processing. A remarkable observation made by R. Coifman and Y. Meyer at the International Congress of Mathematicians in Kyoto, Japan allows us to construct smooth orthogonal bases subordinate to arbitrary partitions of the line. The bases consist of sines or cosines multiplied by smooth cutoff functions. The resulting libraries of localized trigonometric functions can be combined into a tree structure. We will see that this structure is very similar to that of wavelet packets. In particular, we will be able to find a best basis.

Continuous local trigonometric functions. Let $\{I_k\}$ be a collection of disjoint compact intervals of $\mathbf{R}$, indexed by the integers k. Coifman and Meyer showed in [CM2] how to construct an orthonormal basis of $L^2(\mathbf{R})$ subordinate to the partition $\mathbf{R} = \cup_k I_k$ in which the basis vectors are cosines multiplied by smooth bumps.

Let $\ldots c_{-1} < c_0 < c_1 < c_2 < \ldots$ be any increasing sequence such that $c_k \to \infty$ and $c_{-k} \to -\infty$ as $k \to \infty$. We show how to construct an orthonormal basis for L^2 consisting of smooth functions supported on intervals $[c_{k-1}, c_{k+1}]$.

Let $\beta(t)$ be a continuous real-valued function on $\mathbf{R}$ with the following properties:

(1) $\beta(t) = 0$ if $t < -1$,
(2) $\beta(t) = 1$ if $t > 1$,
(3) $\beta(t)^2 + \beta(-t)^2 = 1$ for all t.

For each integer k, define $a_k = \frac{1}{2}(c_k + c_{k-1})$ and $\lambda_k = \frac{1}{2}(c_k - c_{k-1})$. The intervals $I_k = [a_k, a_{k+1}]$ with lengths $|I_k| = a_{k+1} - a_k = \lambda_k + \lambda_{k+1}$ correspond to those defined by Coifman and Meyer. Then we can define the cutoff function b_k for the basis elements supported on $[c_{k-1}, c_{k+1}]$ as follows:

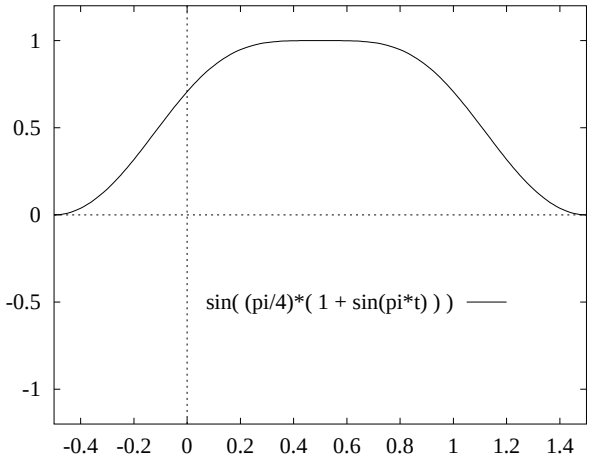
$$b_k(t) = \begin{cases} \beta\left(\frac{t-a_k}{\lambda_k}\right), & \text{if } t \in [c_{k-1}, c_k[, \\ \beta\left(\frac{a_{k+1}-t}{\lambda_{k+1}}\right), & \text{if } t \in [c_k, c_{k+1}]. \end{cases} \tag{5.1}$$

For definiteness we can use a particular bump function

$$\beta(x) = \begin{cases} \sin\frac{\pi}{4}(1+\sin\frac{\pi}{2}t), & \text{if } -1 \le t \le 1, \\ 0, & \text{if } t < -1, \\ 1, & \text{if } t > 1. \end{cases} \tag{5.2}$$

We can plot the cutoff centered at $t = \frac{1}{2}$, and supported on $[-\frac{1}{2}, \frac{3}{2}]$ to illustrate the construction:

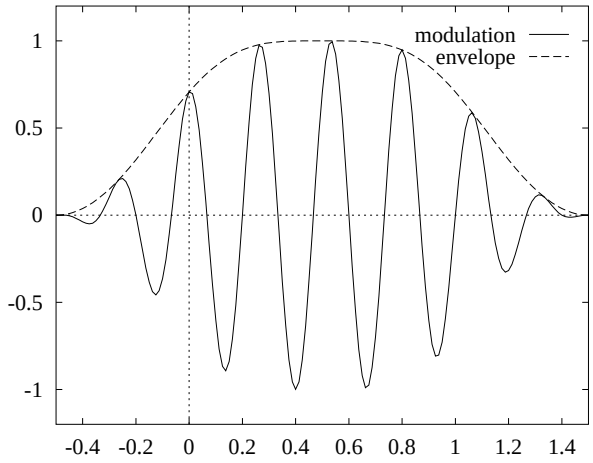
Orthogonal bell for local trigonometric transforms



Since β is smooth on $(-1, 1)$ with vanishing derivatives at the boundary points, b has a continuous derivative on $\mathbf{R}$. Notice that we can modify β to obtain more continuous derivatives by iterating the innermost sine: n iterations will yield at least 2^{n-1} vanishing derivatives at -1 and 1.

Let $m(n,t) = \sqrt{2}\cos\pi nt$, and consider the set of functions $m(n+\frac{1}{2}, t)$ for $n \ge 0$, $t \in \mathbf{R}$. These form an orthonormal basis for $L^2([0,1])$. They may also be dilated and translated to the interval I_k by the formula $m_k(n,t) = \frac{1}{\sqrt{|I_k|}} m(n, \frac{t-a_k}{|I_k|})$. Now define $\psi_{nk}(t) = b_k(t) m_k(n+\frac{1}{2}, t)$ for integers $n \ge 0$ and k, one of which is plotted below.

Local cosine transform basis function



The collection $\{\psi_{nk} : n \geq 0; k \in \mathbf{Z}\}$ forms an orthonormal basis for $L^2(\mathbf{R})$: the proof is a direct calculation. Cosine may be replaced by sine, and there are some other modifications possible, but this set of functions is sufficiently general for our present purposes.

Observe that ψ_{nk} is well localized in both space and frequency. In space, it is compactly supported on $[c_{k-1}, c_{k+1}]$. In frequency, $\hat{\psi}_{nk}$ consists of two modulated bumps centered at $n + \frac{1}{2}$ and $-n - \frac{1}{2}$, respectively, with spread equal to that of $\hat{b}_k$. But $\hat{b}_k$ is well localized because b_k is smooth. We will calculate the variance of $\hat{b}_k$ below.

If we define the functions $\Psi_{nk}(x) = \psi_{n_1 k_1}(x_1) \dots \psi_{n_d k_d}(x_d)$ for multi-indices n and k, we will obtain an orthonormal basis for $\mathbf{R}^d$ made of tensor products. Of course, it is possible to use a different partition in each dimension, as well as different windows.

Discrete local cosine transforms. The functions ψ_{nk} have discrete analogues which form a basis of $l^2(\mathbf{Z})$, or $l^2(\mathbf{T})$. For the former, let $\{c_k : k \in \mathbf{Z}\}$ be an increasing sequence of integers, and define the discrete functions $\tilde{b}_k(j)$ and $\tilde{\psi}_{nk}(j)$ as follows:

$$
\begin{aligned}
\tilde{b}_k(j) &= b_k(j + \frac{1}{2}), \qquad \text{if } c_{k-1} \leq j < c_{k+1}, \\
\tilde{m}_k(n, j) &= m_k(n + \frac{1}{2}, j + \frac{1}{2}), \qquad \text{if } c_{k-1} \leq j < c_{k+1} \text{ and } 0 \leq n < \frac{1}{2}(c_{k+1} - c_{k-1}), \\
\tilde{\psi}_{nk}(j) &= \tilde{b}_k(j)\, \tilde{m}_k(n, j), \qquad \text{if } c_{k-1} \leq j < c_{k+1} \text{ and } 0 \leq n < \frac{1}{2}(c_{k+1} - c_{k-1}).
\end{aligned}
\tag{5.3}
$$

For each k, the functions $\tilde{m}_k$ are evidently the basis functions for the so-called DCT-IV transform [RY] on the interval I_k. The particular bells chosen allow cosines on adjacent intervals to overlap while remaining orthogonal. A similar basis may be constructed over equispaced points on the circle $\mathbf{T}$.

Implementation by folding. Rather than calculate inner products with the sequences ψ_{nk}, we can preprocess data so that standard fast DCT-IV algorithms may be used. This may be visualized as "folding" the overlapping parts of the bells back into the interval. This folding can be transposed onto the data, and the result will be disjoint intervals of samples which can be "unfolded" to produce smooth overlapping segments.

Suppose we wish to fold a function across 0, onto the intervals $[-1/2, 0)$ and $(0, 1/2]$, using the bell b defined above. Then folding replaces the function $f = f(t)$ with the left and right parts f_- and f_+:

$$\begin{aligned} f_-(t) &= b(-t)f(t) - b(t)f(-t), \qquad \text{if } t \in [-\tfrac{1}{2}, 0), \\ f_+(t) &= b(t)f(t) + b(-t)f(-t), \qquad \text{if } t \in (0, \tfrac{1}{2}]. \end{aligned} \tag{5.4}$$

The symmetry of b allows us to use $b(-t)$ instead of introducing the bell attached to the left interval.

Unfolding reconstructs f from f_- and f_+ by the following formulas:

$$f(t) = \begin{cases} b(t)f_+(-t) + b(-t)f_-(t), & \text{if } t \in [-\tfrac{1}{2}, 0), \\ b(t)f_+(t) - b(-t)f_-(-t), & \text{if } t \in (0, \tfrac{1}{2}]. \end{cases} \tag{5.5}$$

Composing these relations yields $f(t) = \big(b(t)^2 + b(-t)^2\big)f(t)$, which is verified by the bell b defined above, for which the sum of the squares is 1. We can translate and dilate these relations to all adjacent pairs of intervals, and of course it works for sequences as well.

General orthogonal bells. Let $b = b(t)$ be any bell. Then $\frac{d}{dt}[b(t)]^2$ is a pair of symmetric continuous bump functions supported respectively on $[c_{-1}, c_0]$ and $[c_0, c_1]$, where c_0 is the center point of the bell, and c_{-1} and c_1 are the center points of the previous and next bells, respectively. This observation provides a mechanism for parametrizing bell functions.

Frequency localization. How well localized the basis functions are in frequency depends upon the relative steepness of the sides of the bell. Suppose first that the bell is symmetric about $t = 0$ and supported in the interval $[-a, a]$, so that a may be called the characteristic width of b. There is a real-valued function $\eta(t)$ supported on $[-a/2, a/2]$ such that $\int \eta = 1$ and

$$b(t) = \int_{-\infty}^{t} [\eta(u - \frac{a}{2}) - \eta(\frac{a}{2} - u)]\, du. \tag{5.6}$$

But then

$$\hat{b}(\xi) = \frac{1}{i\xi}[e^{ia\xi/2}\hat{\eta}(\xi) - e^{-ia\xi/2}\hat{\eta}(-\xi)] \tag{5.7}$$

We also know that $\int |b(t)|^2\,dt = a$, and so $\frac{1}{a}|\hat{b}(\xi)|^2$ really is a probability density function. We can compute the variance of that density as follows:

$$\begin{aligned}\frac{1}{a}\int \xi^2|\hat{b}(\xi)|^2\,d\xi &= \frac{1}{a}\int |e^{ia\xi/2}\hat{\eta}(\xi) - e^{-ia\xi/2}\hat{\eta}(-\xi)|^2\,d\xi \\ &= \frac{1}{a}\int |\hat{\eta}(\xi)|^2\,d\xi + \frac{1}{a}\int |\hat{\eta}(-\xi)|^2\,d\xi - \frac{2}{a}\int \Re\{e^{ia\xi}\hat{\eta}(\xi)\overline{\hat{\eta}(-\xi)}\}\,d\xi \\ &= \frac{2\pi}{a}\int |\eta|^2 - \frac{2}{a}\Re\{\int e^{ia\xi}\hat{\eta}^2(\xi)\,d\xi\} \\ &= \frac{2\pi}{a}\int [\eta^2(t) - \eta(t)\eta(a-t)]\,dt\end{aligned} \tag{5.8}$$

Minimizing this variance is an exercise in the valculus of variations. Alternatively, we can set up the finite-dimensional optimization problem for the discrete bell and its Fourier transform.

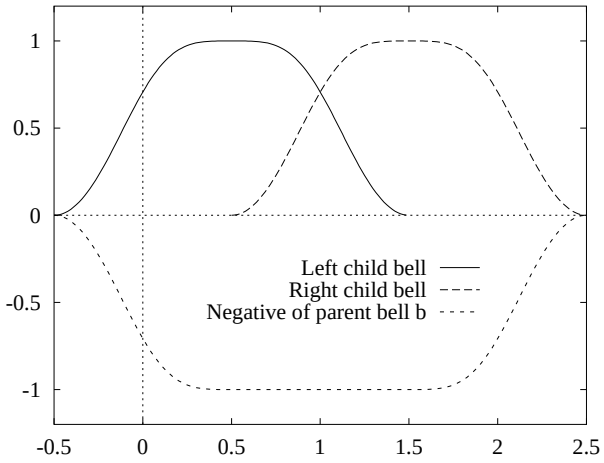
6. Adapted Local Trigonometric Transforms

The subspace of $L^2(\mathbf{R})$ spanned by the local trigonometric functions in adjacent windows are orthogonal, and their direct sum equals the space of local trigonometric functions in the parent window. In the notation of the previous chapter, we have

$$b(t)^2 = b_k(t)^2 + b_{k+1}(t)^2, \tag{6.1}$$

where b_k and b_{k+1} are bell functions supported on $[c_{k-1}, c_{k+1}]$ and $[c_k, c_{k+2}]$, respectively. Thus b is supported on $[c_{k-1}, c_{k+2}]$. The overlap of these bells is displayed below:

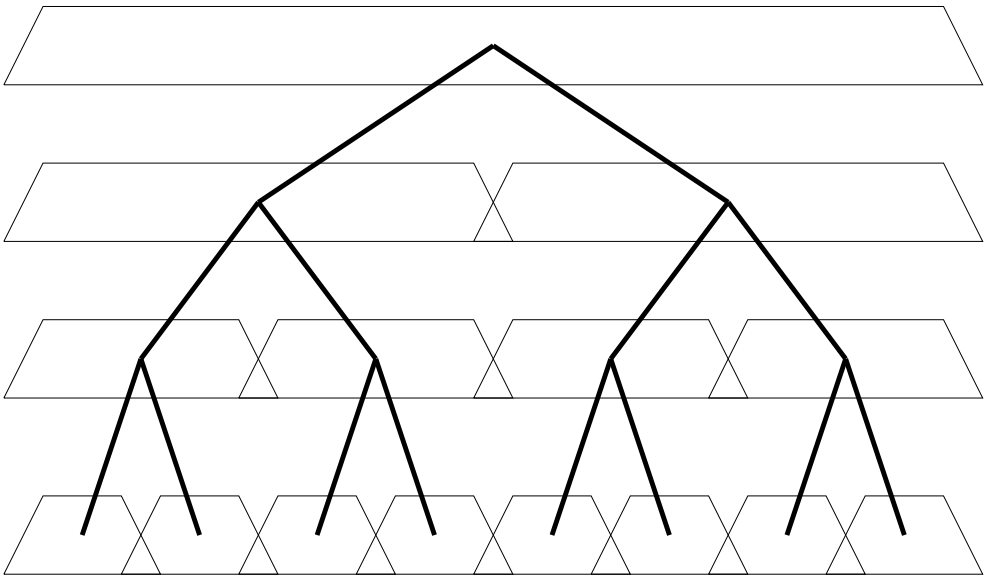
Combining adjacent windows



Adjacent parent windows span orthogonal subspaces too, which may be combined into grandparent subspaces, and so on. The process recursively builds a tree of subspaces with the same remarkable orthogonality properties as the wavelet packet tree, only in this library each space contains the frequencies at a fixed location. This is like the transpose of the wavelet packet tree.

We may draw the tree of the local trigonometric function library schematically as below, where we restrict ourselves to just a few levels.

The tree of adjacent orthogonal windows

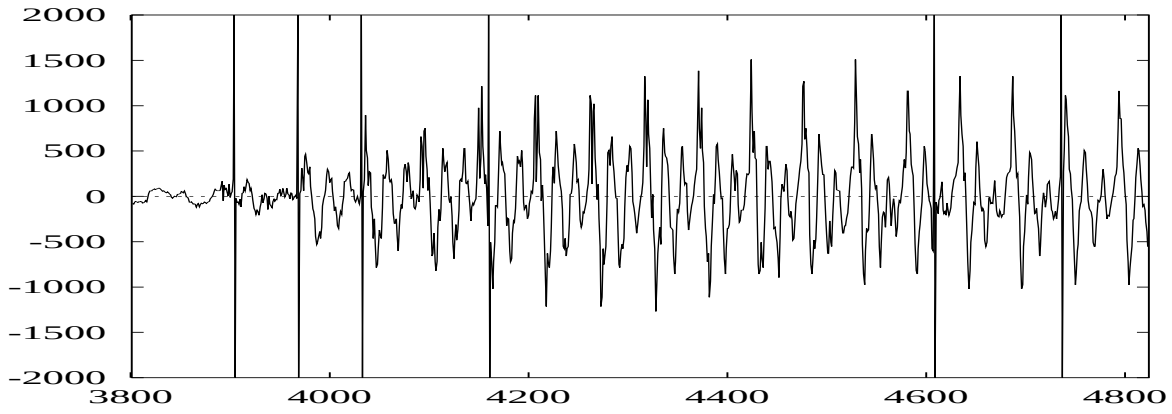


Since the local trigonometric transform imposes no restrictions on the support intervals of the bells, the subspaces need not be of equal size. Nor is it necessary to combine windows in pairs: the tree can be inhomogeneous if desired. We need not concern ourselves with these evident generalizations.

Signal segmentation. The best basis algorithm applies to this tree, and has the following interpretation. If the frequencies present in two adjacent windows are sufficiently similar, so that their parent window has a lower information cost, then represent the signal in that region in the parent space. Apply this recursively from the bottom of the tree up.

This algorithm segments the signal by efficiency of representation. It chooses small windows to cover transients with rapidly changing frequency content (such as consonants), and large windows to cover long regions with relatively constant frequency content (such as vowels). Below is a graph prepared by Xiang Fang, depicting the window sizes chosen in a small portion of the spoken word "armadillo." The vertical spikes were introduced artificially to indicate where the algorithm chose the window boundary points a_k.

$\frac{1}{8}$ *second of "armadillo, " near the "d"*



The binary tree structure introduces artifacts because nodes which are close together in time are not necessarily close together in the tree. As a result, certain adjacent pairs of segments which may be virtually identical in frequency content may not be combined by the algorithm. Also, the initial placement of the narrowest window boundary affects the choice of segments. It can be shown that the number of segments chosen varies at most by a factor equal to the number of levels.

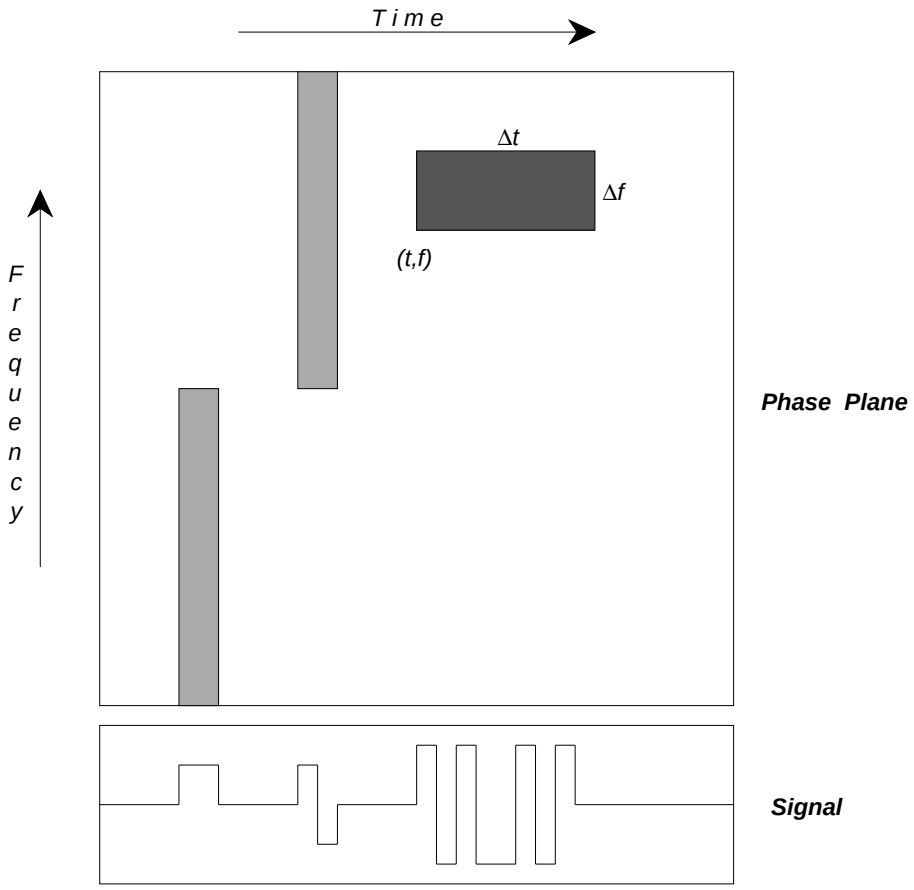
Another difficulty is that the edge of the widest bell is exactly as steep as that of the narrowest. This introduces inefficiency if there are many levels, since the entropy of

the cosine transform coefficients within a smooth bell is controlled by the product of the steepness of the bell and the width of the window. Ideally, this product should be held nearly constant for all window widths. There is a way to make this ratio constant, by sacrificing some of the time localization, but we will not consider it in these notes.

7. Time-Frequency Analysis

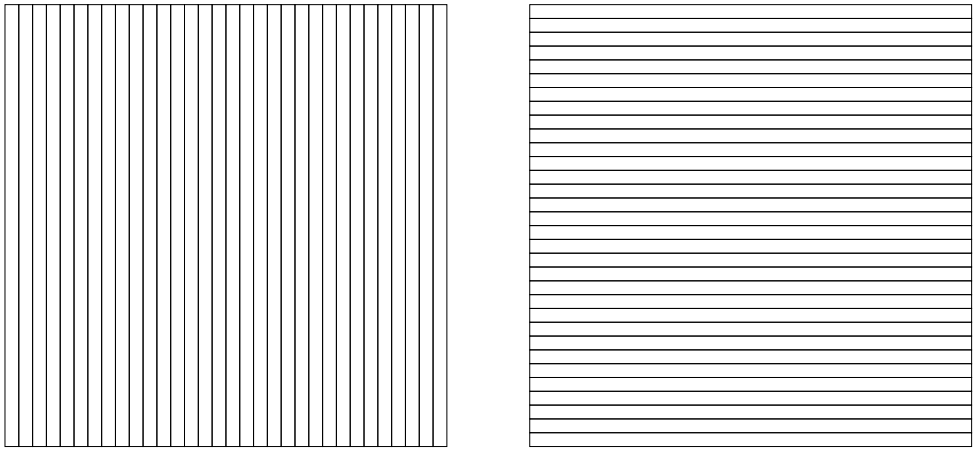
To each wavelet packet or local trigonometric function we can associate a time t and a frequency f. These will be uncertain by amounts Δt and Δf, respectively. The result may be interpreted as a rectangular patch of dimensions Δt by Δf, located around (t, f). We shall call the patch a *Heisenberg box* in honor of the uncertainty principle, which limits how small the area of the patch may be. The Heisenberg box may be colored in proportion to the amplitude of the corresponding wavelet packet component.

Heisenberg Boxes in the Phase Plane.



An orthonormal basis corresponds to a disjoint cover of the phase plane by Heisenberg boxes. Certain bases have characterizations in terms of the shapes of the boxes present in the cover. For example, the standard basis consists of the cover by the tallest, thinnest patches allowed by the sampling interval. The Fourier transform may be regarded as the transpose of the standard basis, in the sense that the Heisenberg boxes are transposed by interchanging time and frequency. The standard basis has optimal time localization and no frequency localization, while the Fourier basis has optimal frequency localization, but no time localization.

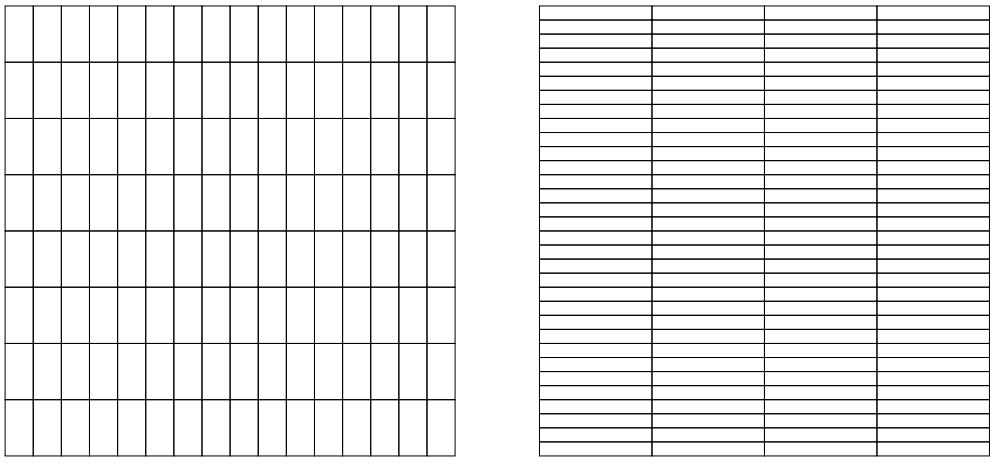
Phase Plane Decomposition by the Standard and Fourier Bases.



Standard Basis Fourier Basis

Windowed Fourier or cosine transforms with a fixed window size correspond to covers with congruent Heisenberg boxes whose width Δt is the window width. The ratio of frequency uncertainty to time uncertainty is the aspect ratio of the Heisenberg boxes:

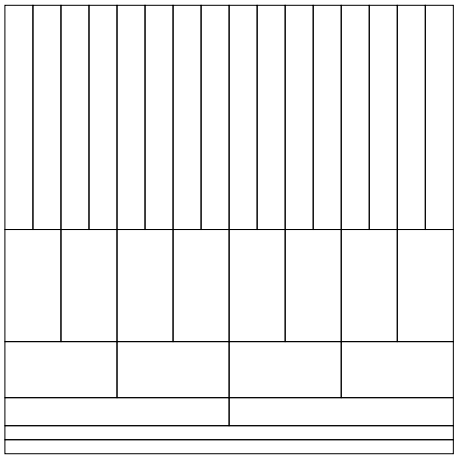
Phase Plane Decomposition by Windowed Cosine Transforms.



Narrow windows Wide windows

The wavelet basis is an octave-band decomposition of the phase plane, depicted by the following cover:

Phase Plane Decomposition by Wavelet Transform.



Wavelet basis

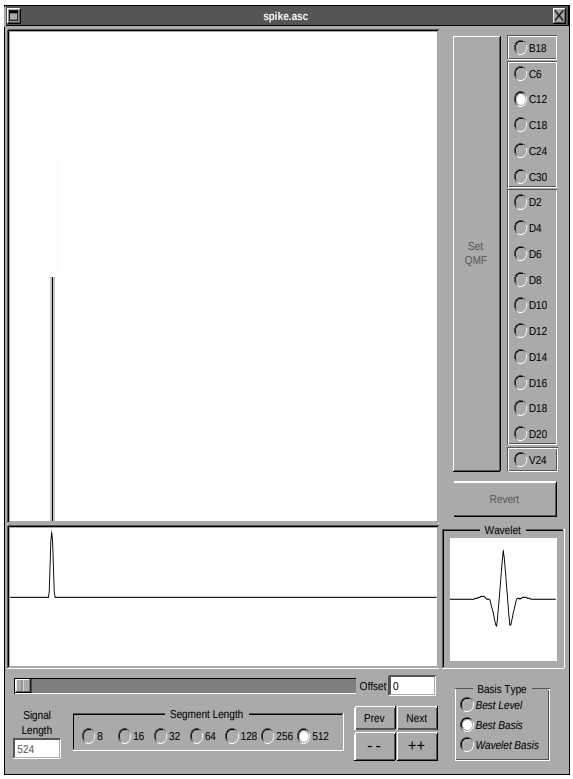
The best-basis of wavelet packets fits a cover to the signal so as to minimize the amount of dark Heisenberg boxes. The compressibility of a sampled signal is easily seen to be the ratio of the total area of the phase plane ($N \times N$ for a signal sampled at N points) divided by the total area of the dark Heisenberg boxes (each of area N). This method allows rectangles of all aspect ratios. The best-level or adapted subband basis fits a cover of equal aspect ratio rectangles to the signal, so as to minimize the amount of dark.

We may automatically analyze signals by expanding them in the best basis, then drawing the corresponding phase plane representation. As is clear, the negligible components will not be drawn, as it is not relevant which particular basis is chosen for a subspace containing negligible energy.

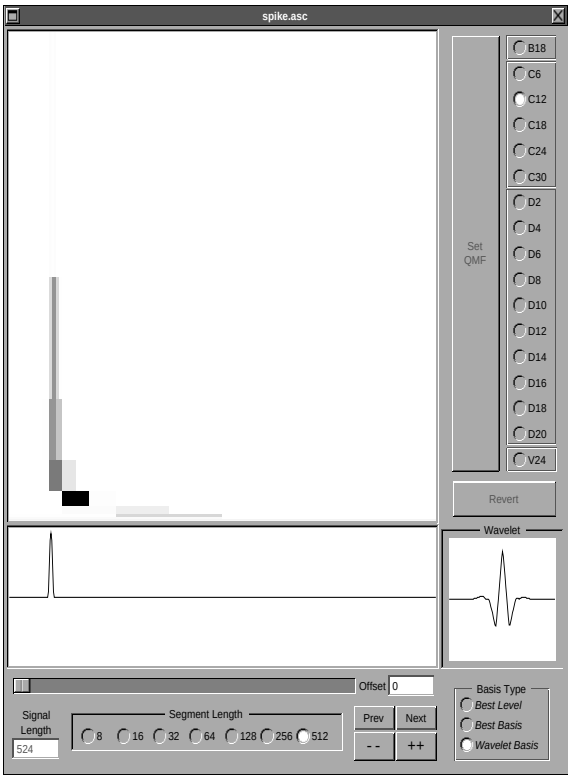
Below are certain canonical signals and their automatic analyses by a wavelet packet program written for a desktop computer. The user selects a quadrature mirror filter from a list of 17 at the right, and the "mother wavelet" determined by that filter is displayed in the small square window at the lower right. The signal is plotted in the rectangular window at bottom, and the phase plane representation is drawn in the large main square window. Instructions for obtaining and using the software package used in this paper may be found in Appendix 1.

We first analyze a relatively smooth transient, spread over 7 samples in a 512 sample signal:

Representing a Fast Transient



Best basis

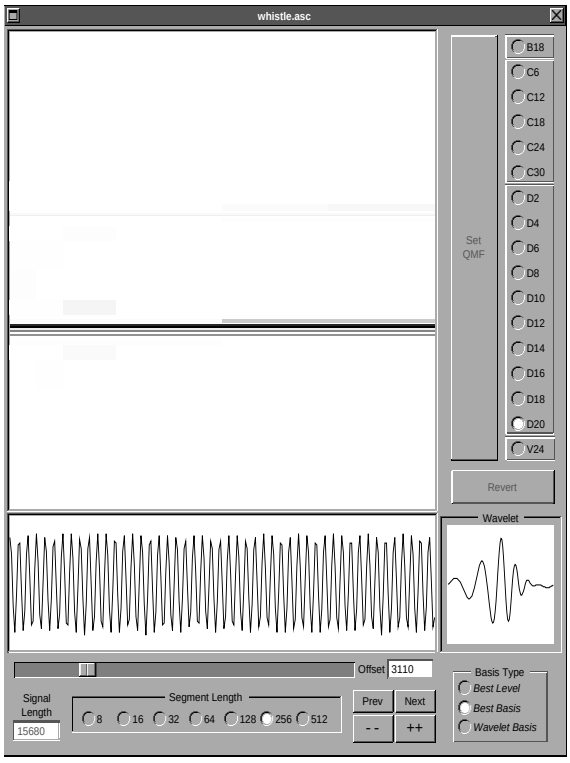


Wavelet basis

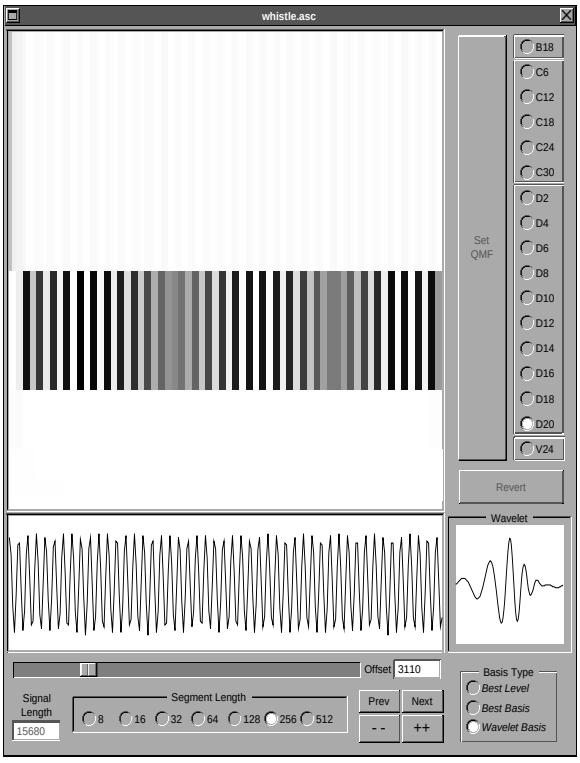
Notice that the wavelet analysis at the right correctly localizes the peak in the high-frequency components, but is forced to include poorly localized low-frequency elements as well. The best-basis analysis finds the optimal representation within the library, which in this case is almost a single wavelet packet.

The second signal is taken from a recording (at 8012 samples per second) of a person whistling:

Representing a Whistle



Best basis

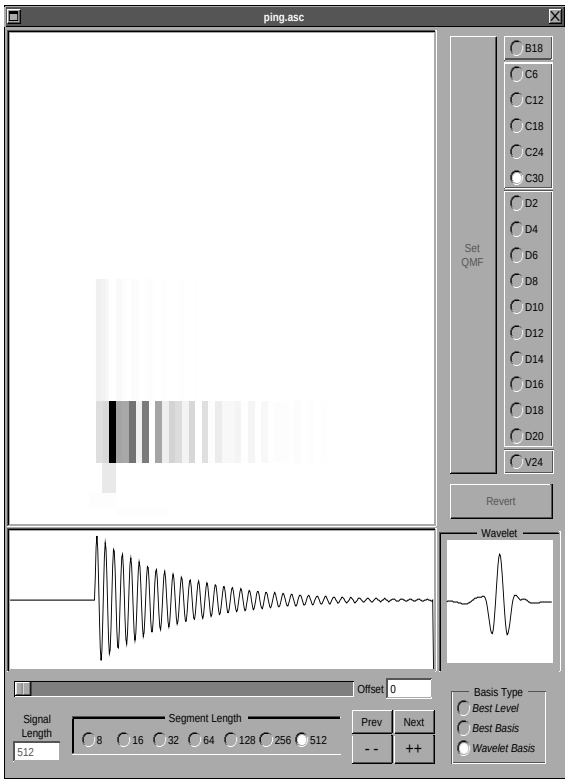


Wavelet basis

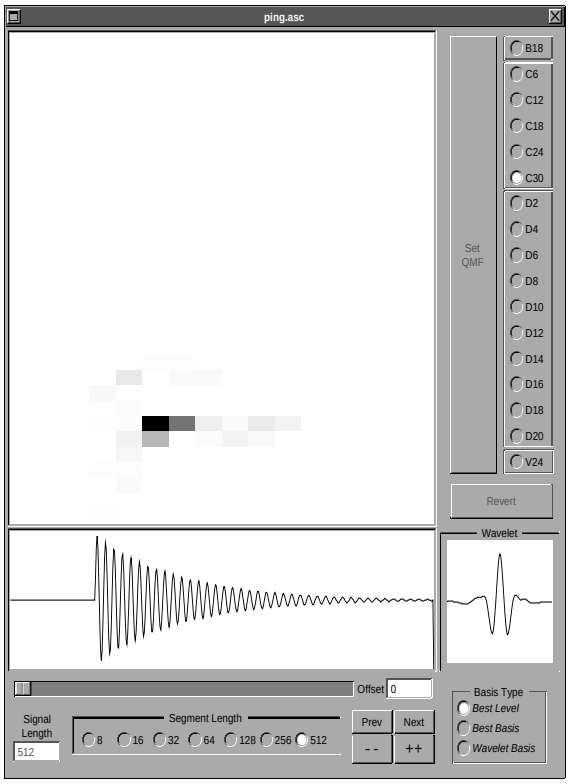
Here the wavelet basis is only able to localize the frequency within an octave, even though the best-basis analysis shows that it falls in a much narrower band. The vertical stripes among the wavelet Heisenberg boxes may be used to further localize the frequencies, but the best-basis decomposition performs this analysis automatically.

Let us now combine the transient and periodic parts in different ways. For example, we may take a critically damped oscillator which receives an impulse, and decompose the resulting solution in the wavelet and best-level bases, as below. The wavelet decomposition locates the discontinuity at the impulse, while the best-level analysis finds the resonant frequency of the oscillator more precisely.

Critically Damped Harmonic Oscillator



Wavelet basis

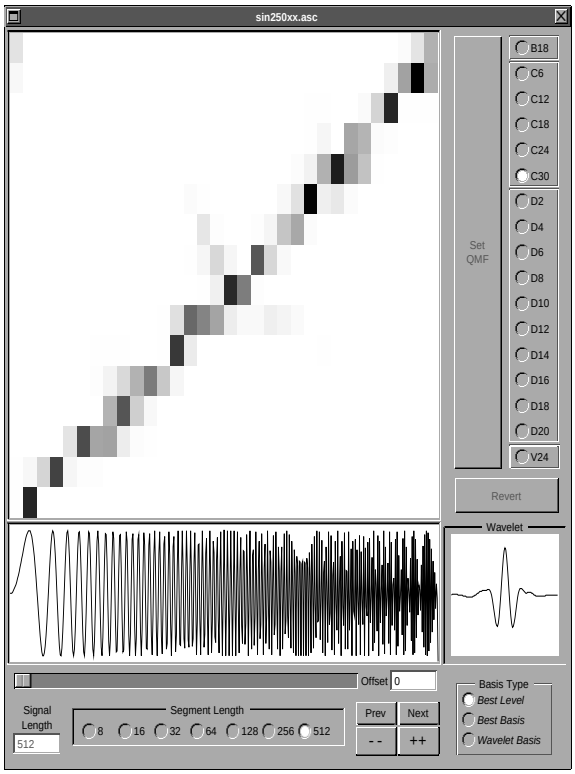


Best-level basis

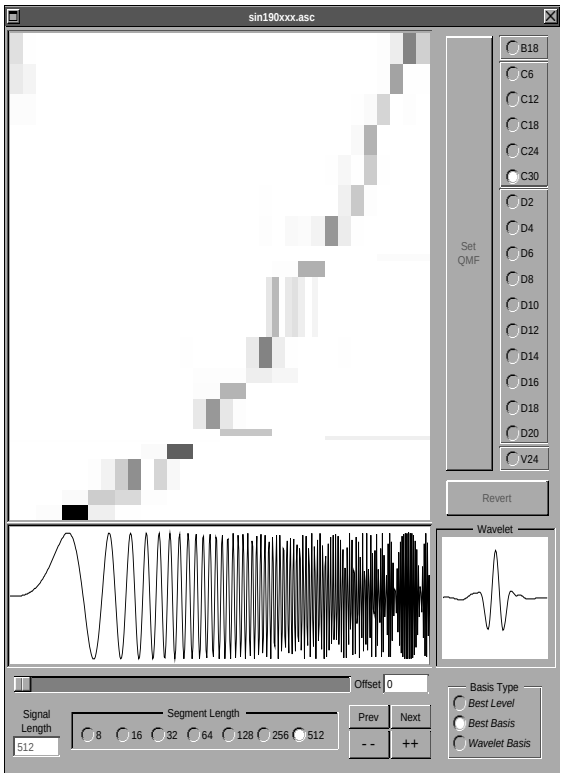
The exponential decay of the amplitude is visible in both analyses.

A chirp is an oscillatory signal with increasing modulation. For example, below are the functions $\sin(250\pi t^2)$ and $\sin(190\pi t^3)$ on the interval $0 < x < 1$, sampled 512 times. The modulation increases linearly and quadratically, respectively. The Heisenberg boxes form a line and a parabolic arc, respectively. In the best-level analyses, all the Heisenberg boxes have the same aspect ratio, which is appropriate for a line. In the best-basis analysis, the Heisenberg boxes near the zero-slope portion have smaller aspect ratio than those near the large-slope portion.

Linear and Quadratic Chirps



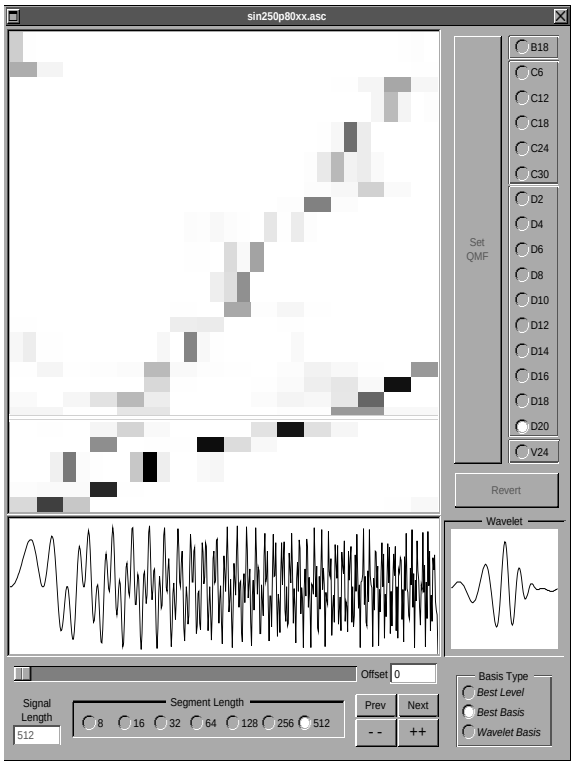
Linear chirp, best level



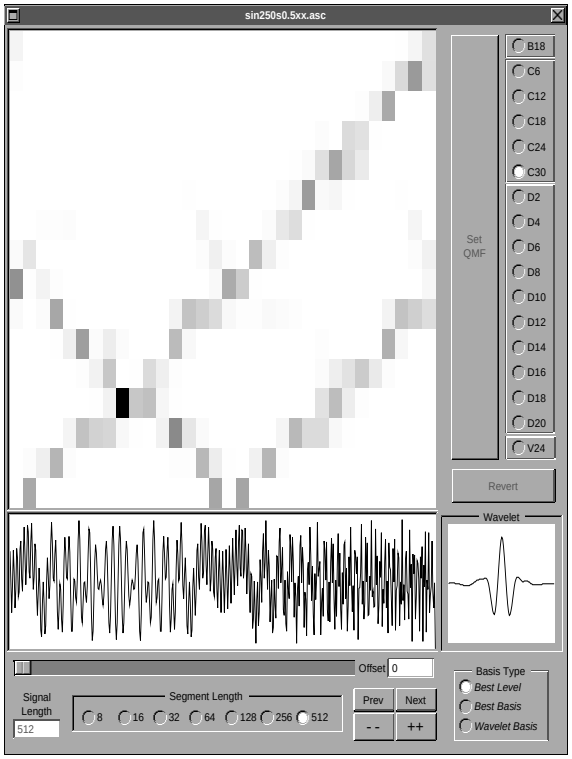
Quadratic chirp, best basis

Such a time-frequency analysis can separate superposed chirps. Below are pairs of linear chirps, differing either by modulation law or phase. Both are functions on the interval $0 < t < 1$, sampled 512 times. On the left is the function $\sin(250\pi t^2) + \sin(80\pi t^2)$ analyzed in the best wavelet packet basis. Note that the milder slope chirp is represented by Heisenberg boxes of lower aspect ratio. On the right is $\sin(250\pi t^2) + \sin(250\pi(t - \frac{1}{2})^2)$, analyzed by best-level wavelet packets. The downward-sloping line comes from the aliasing of neagative frequencies.

Superposed Chirps



Different slopes, best basis



Different phases, best level

8. Multimensional Wavelet Packets

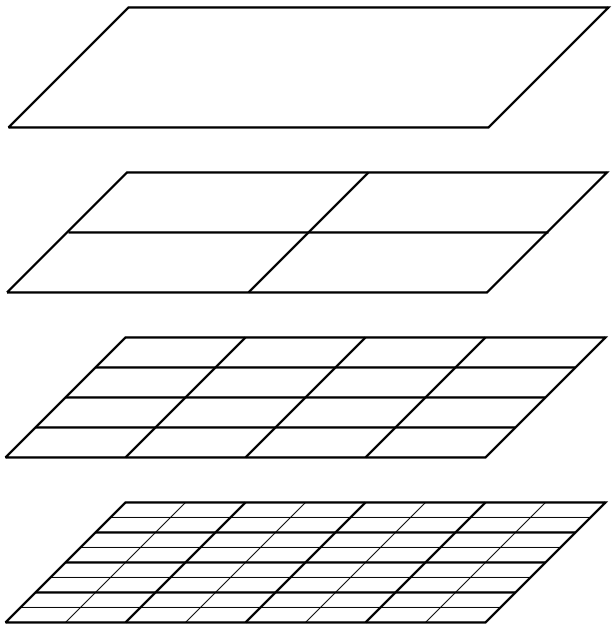
The quadrature mirror filter algorithm extends to multidimensional signals by separation of variables. Let F_0 and F_1 be a pair of QMF's for $\mathbf{R}$; then $F_0 \otimes F_0$, $F_0 \otimes F_1$, $F_1 \otimes F_0$, and $F_1 \otimes F_1$ are a family of 4 filters for $\mathbf{R}^2$. In general, for d-dimensional signals, we will use a 2^d member family of filters $\otimes_{i=1}^{d} F_{\epsilon_i}$, where $\epsilon_i \in \{0, 1\}$ for $i = 1, \ldots, d$. The wavelet packets produced by iterating these filters are products of one-dimensional wavelet packets $W_n(x) = \prod_{i=1}^{d} W_{n_i}(x_i)$, together with their isotropic dilations and translations to arbitrary lattice points.

Inner products with these multidimensional wavelet packets are computed from averasges at the smallest scale, just as in the one-dimensional case. The coefficients may be organized into a stack of d-dimensional intervals, and there is a result analogous to

Theorem 1.2 which characterizes orthogonal basis subsets.

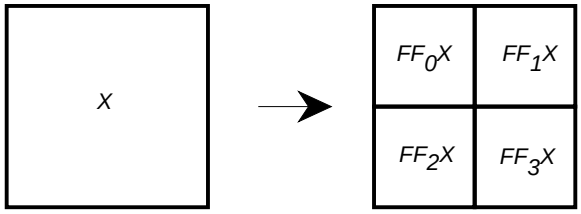
We can restrict our attention to the case $d = 2$. The diagram below depicts a decomposition of the unit square into a 3-level wavelet packet tree:

2-dimensional wavelet decomposition to level 3



Abusing notation, we can write $FF_0 = F_0 \otimes F_0$, $FF_1 = F_0 \otimes F_1$, $FF_2 = F_1 \otimes F_0$, and $FF_3 = F_1 \otimes F_1$. Note that the numbering corresponds to $\epsilon_1\epsilon_2$ in binary. Each 2-dimensional filter decimates by 2 in both in the x and y directions, so it reduces the number of coefficients by 4. The coefficients in the various boxes of this picture are calculated by the application of the filters as labelled below.

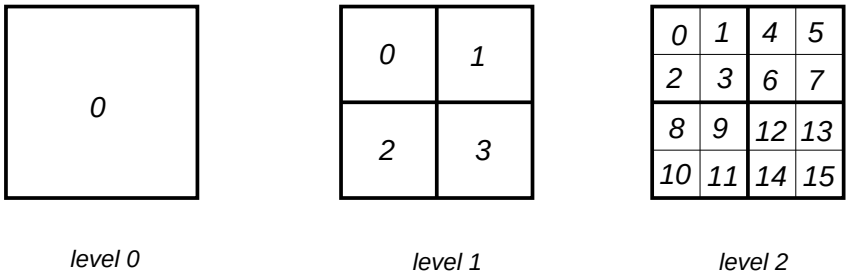
2-dimensional wavelet packet coefficients



By applying the filter convolutions recursively, we obtain a homogeneous tree-structured

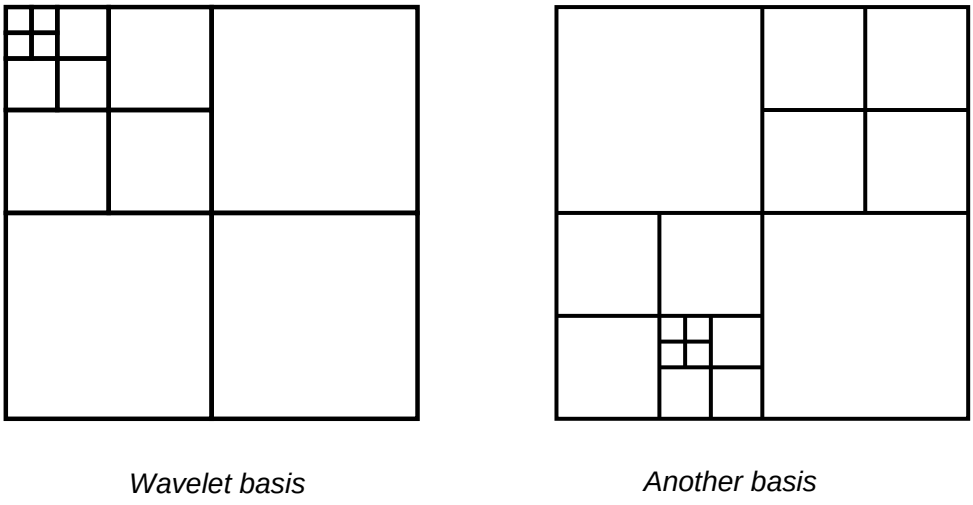
decomposition organized by frequency. We may label the various subblocks as below:

Numbering of 2-dimensional wavelet packets



A basis in the d-dimensional case is a disjoint cover of the unit d-cube with dyadic d-cubes. In the 2-dimensional case, we can draw examples of such covers by squares as below:

Example 2-dimensional wavelet packet bases



Complexity of the d-dimensional best-basis algorithm. Suppose for simplicity that we start with a d-dimensional periodic signal of size $N = 2^L \times \cdots \times 2^L = 2^{Ld}$. This can be developed down to level L, at which point each subspace will have a single element. Each level requires $O(N)$ operations, where the constant is proportional to the product of d and the length of the quadrature mirror filters. The total complexity for calculating the wavelet packet coefficients—i.e., correlating with the entire collection of modulated

bumps—is $O(N \log N)$.

Let A_L be the number of covers in the tableau from levels 0 to L in the d-dimensional case. Then $A_0 = 1$ and $A_{L+1} = 1 + A_L^{2^d}$, so we have an estimate that $A_{L+1} \geq 2^{2^{Ld}}$. In the periodic case with a signal sampled at $N = 2^{Ld}$ points, this implies that there are more than 2^N bases. The tree contains $1 + 2^d + \cdots + (2^d)^L = (2^{(L+1)d} - 1)/(2^d - 1) = O(N)$ subblocks, each of which is examined at most 3 times in the search for a best basis. This shows that the extraction of a best basis has complexity $O(N)$ with small constant.

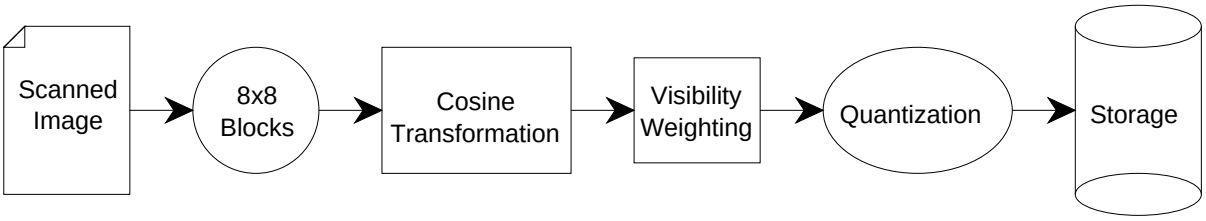
Reconstruction from an arbitrary basis is of the same complexity, since the algorithm is an orthogonal transformation.

9. Picture Compression

A particular use for 2-dimensional transform coding is analysis and compression of pictures.

The JPEG algorithm. If the pixel values in the picture vary smoothly, i.e., are highly correlated with their nearest neighbors, then we can gain efficiency by decorrelating them in another basis. This proceadure is the foundation for classical Karhunen–Loéve based techniques, which treat small sections of the signal as independent random vectors satisfying Gauss-Markov statistics of a particular kind. Under this model, the Karhunen–Loéve basis consists of cosines. The standard transform coding algorithm proposed by the Joint Photographic Experts Group (JPEG) involves segmenting the picture into small squares with 8 pixel sides, then performing a two-dimensional discrete cosine transform within each square. The resulting coefficients fall into two categories: the zero-frequency or "DC" component, and the "AC" components. Since pixel values are taken to be positive, the DC component is always large and positive, but it varies smoothly from block to block; it may quantized to high accuracy, and then adjacent DC values may be efficiently stored as small differences to exploit the slow variation. The AC components are weighted by their "visibility," which is determined through experiments with human subjects, and then quantized to the precision allowed by the desired compression ratio. Schematically, the algorithm is depicted below:

JPEG picture compression

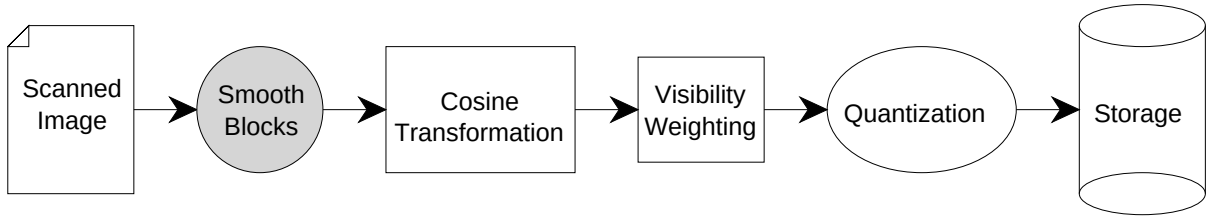


The algorithm and some of its subtleties are described in [RY]. JPEG is remarkably efficient, and various pieces of the algorithm are highly optimized. However, the algorithm can be enhanced in several ways, which will become clear upon considering some of its drawbacks.

Modifications of JPEG. The arbitrary division into 8×8 blocks ignores the large-scale structure of the picture, and prevents the decorrelation of pixels which are not very close together. It also makes the output values less meaningful, since apart from differencing the DC components, no attempt is made to relate values from different blocks to see if they arise from the same large feature. These limitiations can be partially overcome by the best-basis method. The picture may be repartitioned into 16×16 "parent" blocks, one over each quadruplet of 8×8 "child" blocks. The larger blocks of pixels are direct sums of 4 of the smaller blocks. We then compute the information costs of the parent and its children, and see which generation's transformed values are more efficient. Of course, this may be repeated with 32×32 pixel blocks, and so on, up to the full size of the picture. THe best-basis search procedure described in a previous chapter will converge to the optimal segmentation of the picture into dyadic blocks. Repartitioning as many times as possible and searching the (quaternary) tree that results takes $O(N \log N)$ operations, where N is the number of pixels. A parent will be chosen over its descendents if there is a dominant oscillatory component of greater extent than any of the descendents. Significant features of this component are distilled into the size and location of the chosen block, and the frequencies of the large coefficients within it.

The sharp edges of the JPEG blocks pose another problem, since they create artifacts. Part of the engineering effort must be spent filtering out these artifacts, and the effort may not succeed at high compression rates for sensitive data. This difficulty may be overcome by replacing the discrete cosine transform with the local cosine transform, i.e., substituting the shaded step as in the schematic below:

Local cosine picture compression

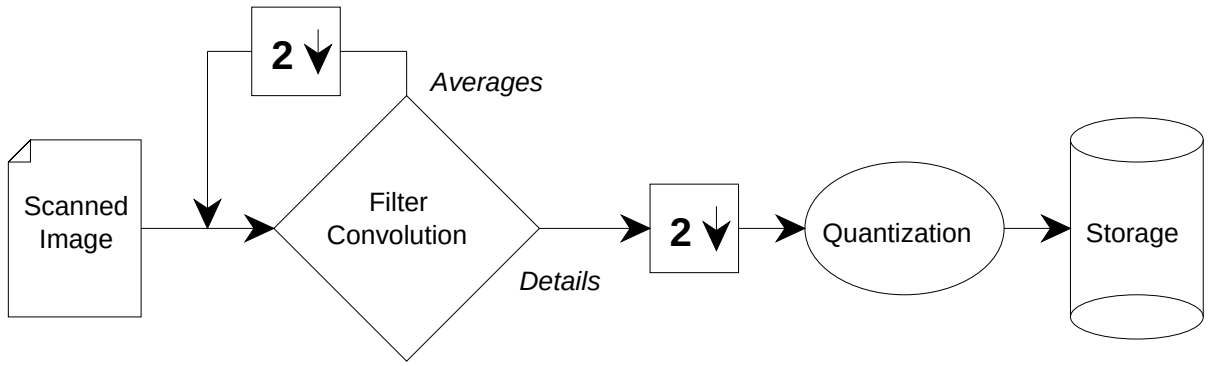


This modified algorithm can also be coupled with an adaptive or best-basis scheme, the two-dimensional version of the adaptive local cosine transform in the chapter above.

Wavelet and wavelet packet methods. JPEG and the variants above perform segmentation in space, either with a fixed model or adaptively to improve efficiency of representation. We may also perform the segmentation by scales, or by both scale and frequency with a subsequent choice of optimal representation.

The wavelet or multiresolution decomposition divides the picture into two pieces, which we may call "averages" and "details." The details correspond to transients at the current scale, while the averages represent transients at larger scales, which look essentially constant at the current scale. The details are stored, and the process is repeated with the averages at the next larger scale. Schematically, the algorithm is depicted below:

Wavelet or multiresolution picture compression

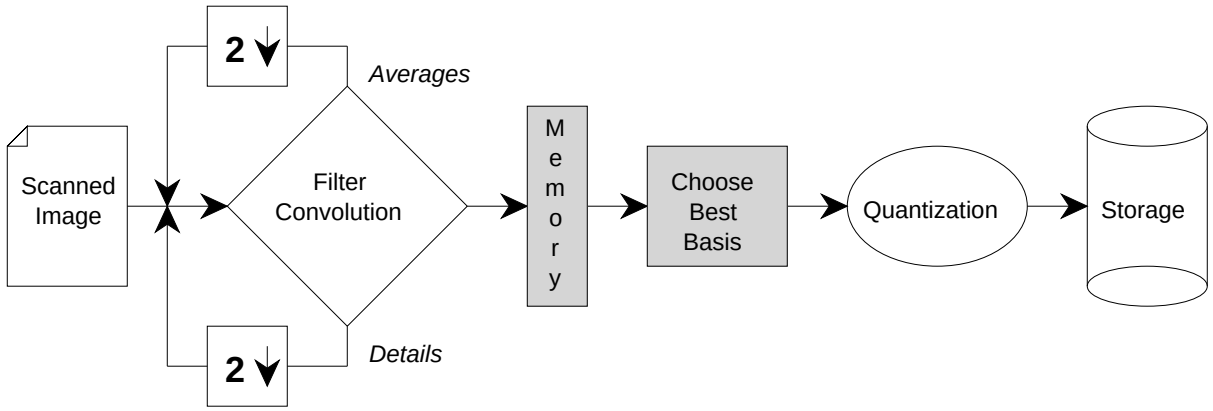


The division into details and averages corresponds to high-pass and low-pass filtering in the quadrature mirror filter algorithm, or to projection onto low-order and high-order

polynomials in an algorithm due to Alpert [Al]. This transformation decorrelates pixel variations by scale, and has proved very sucessful in video compression. It has the advantage that it it segments the picture into layers of increasing detail, so that a picture can be described roughly at first, with more detail provided as needed.

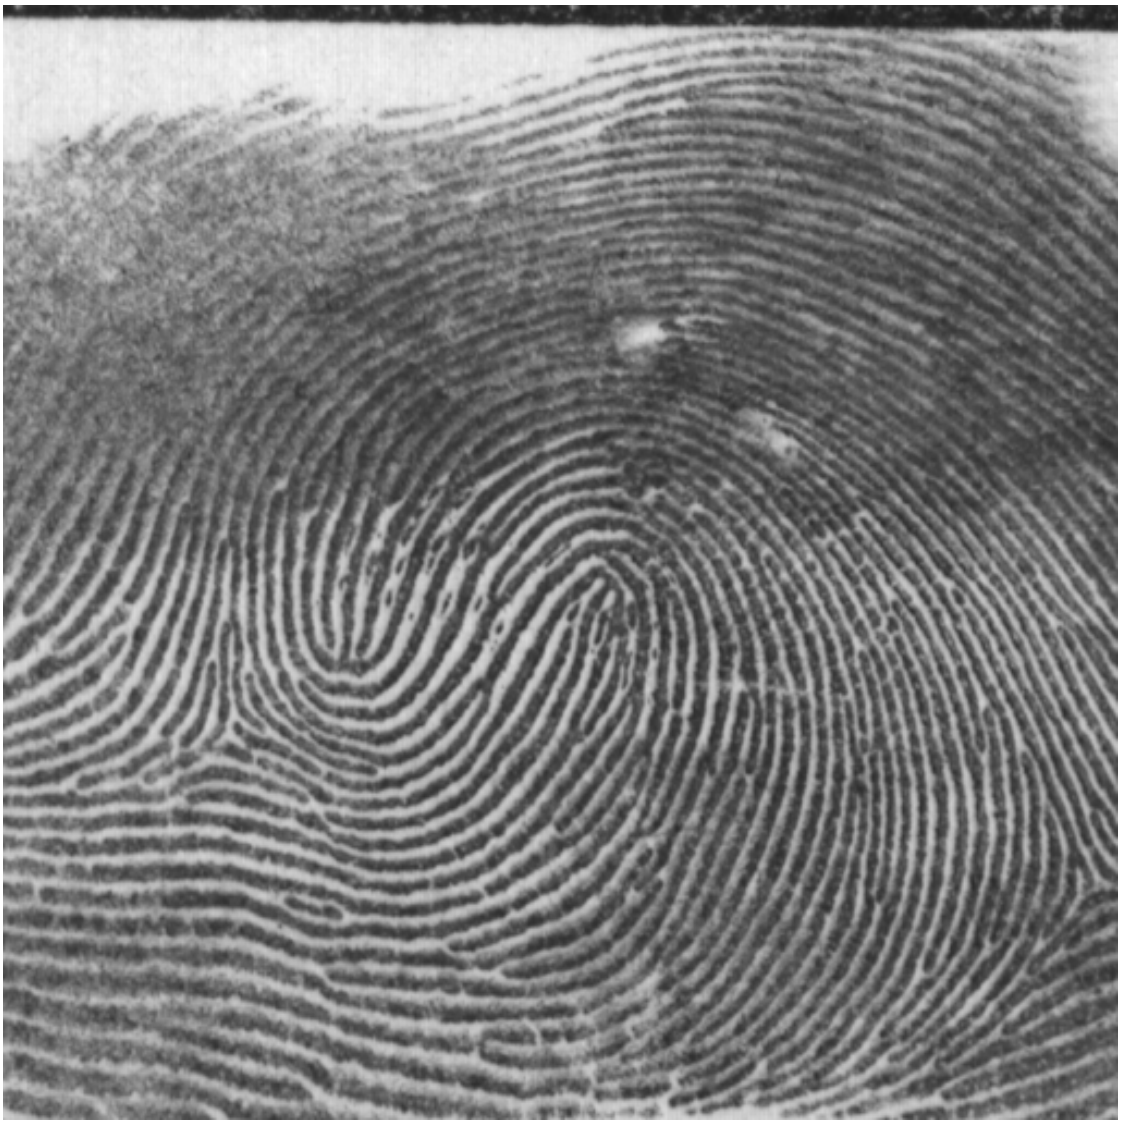
Wavelet packets constitute more than a basis, and provide a finer analysis of the frequency content of transients at a gived scale. This is done by recursively decomposing the details as well as the averages by further filtering. The results may be stored in memory until the possible basis subsets are searched to find the optimal representation. The schematic of this algorithm is given below:

Wavelet packet best-basis picture compression



Comparison of the methods. We shall examine one image to compare these methods. The data was provided by the US Federal Bureau of Investigation for the evaluation of picture compression algorithms for scanned fingerprints, and consists of the central 512×512 pixel portion of an 832×816 image of a thumbprint. The resolution is 500 pixels per inch, 8 bits per pixel.

Original image, 512×512 *pixels, 8 bits per pixel*



(Courtesy of the US Federal Bureau of Investigation)

We now perform a JPEG compression of the image, with the weight matrix given below:

```
{16, 11, 10, 16,  24,  40,  51,  61},
{12, 12, 14, 19,  26,  58,  60,  55},
{14, 13, 16, 24,  40,  57,  69,  56},
{14, 17, 22, 29,  51,  87,  80,  62},
{18, 22, 37, 56,  68, 109, 103,  77},
{24, 35, 55, 64,  81, 104, 113,  92},
{49, 64, 78, 87, 103, 121, 120, 101},
{72, 92, 95, 98, 112, 100, 103,  99}
```

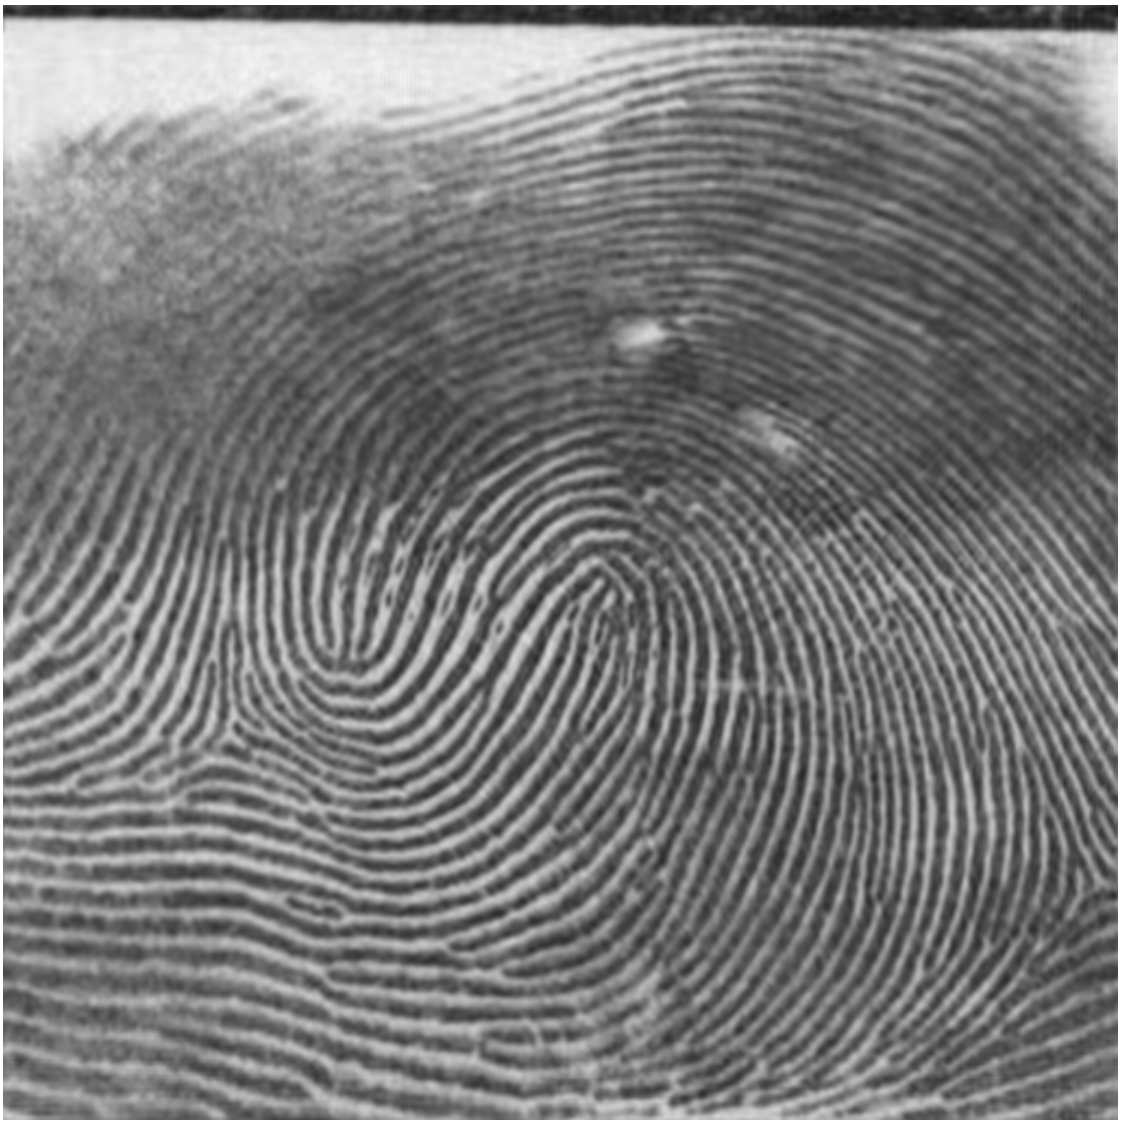
To get a compression ratio of about 20 requires a quantization of 6 bits per coefficient, which introduces some of the blocking effects mentioned earlier:

JPEG picture compression, 6 bits quantization, 0.418bpp, SNR 20.0 dB



A slightly better compression ratio is achieved by local cosine transform with a larger block size, and more precise quantization of the resulting coefficients. This may be thought of as a manual best-level search; the automatic best-basis algorithm in fact selected a 32×32 pixel block size for all but 4 blocks in this example.

Local cosine on 32×32 *blocks, 9 bits quantization, 0.410bpp, SNR 24.1 dB*



A similar compression ratio obtained with the wavelet transform yields a different type of distortion:

Wavelet transform to 9 scales, 9 bits quantization, 0.613 bpp, SNR 15.5 dB



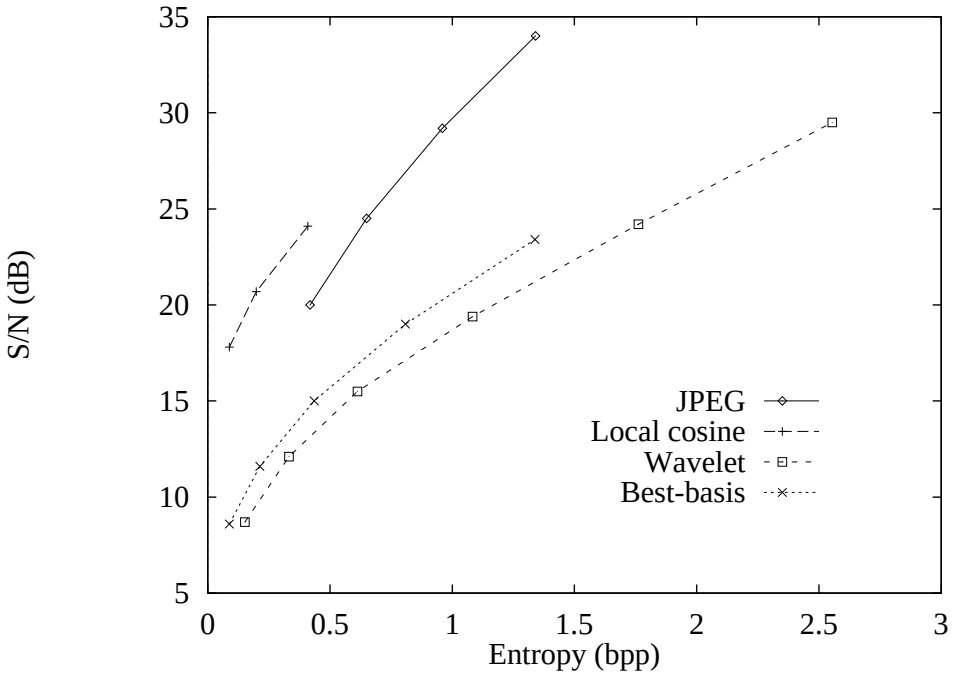
Wavelet packets give a higher compression ratio at the same coefficient quantization:

Wavelet packets to 9 levels, 9 bits quantization, 0.436 bpp, SNR 15.0 dB



To compare these algorithms across a range of compression ratios, we plot the rate-distortion curves for the available quantizations. What is remarkable is that similarly distorted images yield quite dissimilar signal-to-noise ratios.

Rate-distortion curves for four compression algorithms, one picture



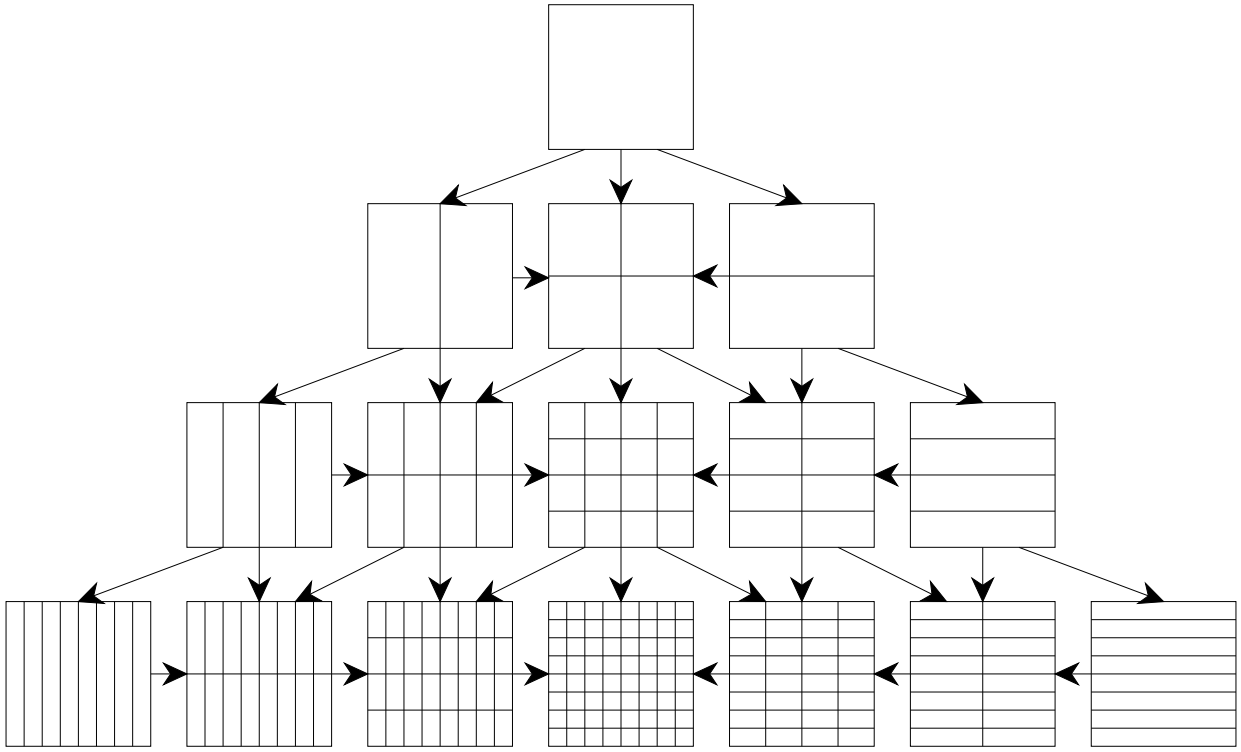
The absence of objectionable artifacts in the wavelet and best-basis compressions, despite their low signal-to noise ratio, speaks for a reëxamination of our error criterion. In fact, the choice of which algorithm to use for such sensitive data must be made on the basis of empirical criteria. The FBI, for example, has designed a series of experiments to evaluate compression methods based upon the time required to obtain a positive identification from compressed fingerprint images.

10. Anisotropic Dilations in Multidimensions

We have so far considered only those multidimensional wavelets which have the same scale in each direction. If we consider bases of wavelet packets whose scales are different in different directions, then the number of coefficients and the number of bases increase substantially. Below is one way to organize the calculation of 2-dimensional wavelet packet coefficients. The arrows represent filter convolution-decimations in one of the dimensions, X or Y. This diagram shows that there will be $O(N[\log N]^2)$ coefficients in the 2-dimensional case. More generally, it can be proved that there will be $O(N[\log N]^d)$ coefficients in the

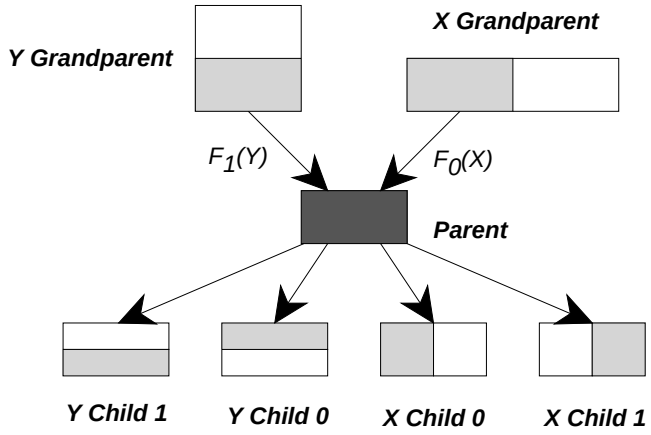
d-dimensional case.

All tensor products of wavelet packets



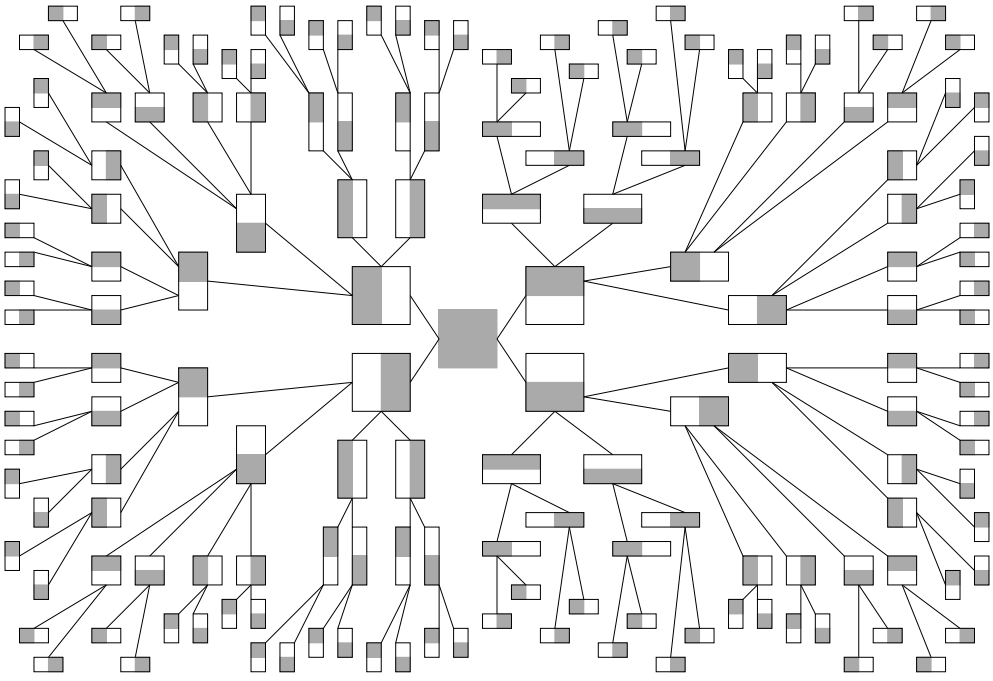
With anisitropic dilations, wavelet packet coefficients no longer form a tree. Each subspace has d parents, one for each filter, as illustrated by the genealogy below:

Genealogy of tensor products of wavelet packets



The resulting structure can be made into an inhomogeneous tree by including duplicate copies of the parent, one for each grandparent, recursively for each generation. The result is a very large tree, as the diagram below for $N = 4 \times 4$ illustrates:

Wavelet packet tree for a 4×4 *signal*



Each node within this tree of dependencies is indexed by the string of convolution-decimation operators that produced it, i.e., $G_m \dots G_2 G_1$, where $m \le 2L$ and where G_i is one of $F_0(X)$, $F_0(Y)$, $F_1(X)$, or $F_1(Y)$. Order is important, and there can be no more than L convolution-decimations in each direction X or Y. Counting these gives

$$\sum_{x=0}^{L} \sum_{y=0}^{L} \frac{(x+y)!}{x!\,y!} 2^x 2^y$$

This can be shown to grow like N^2, where $N = 2^L \times 2^L$ is the number of samples in the picture. In higher dimensions the combinatorial growth is even worse, suggesting that tensor product wavelet packets are an impractical generalization.

11. Compression

Let x be a vector in $\mathbf{R}^N$ with coordinates determined up to some fixed precision b bits. The space required to store x is bN bits. We wish to represent x with fewer bits, either with no losses, or with small errors in an appropriate norm.

Shannon gives a lower bound for the space required to store x with no losses, in terms of the entropy of the distribution of coefficient values. Practical data compression methods such as Huffman coding approach this lower bound for large values of bN. Lossless compression methods exist which apply an invertible transformation $T : x \to x'$ in the expectation that the distribution of coefficients of x' has lower entropy than x. If $\{x_k : k = 1, \dots, K\}$ is an ensemble of vectors produced by a Gaussian process on $\mathbf{R}^N$, then the optimal entropy reduction is obtained when T is the Karhunen–Loéve transformation, which is described below.

Compression by quantization. Further compression is obtained by scalar quantization. Here we replace the coefficients in the transformed signal $x' = Tx$ by b'-bit approximations, where $b' < b$, then compress these by Huffman coding. The last step is nonlinear, but the first step is linear to the precision of the quantization. The quality of such a compression is depicted as a rate-distortion curve parametrized by b'. The rate-distortion curve itself relates the signal to noise energy ratio in decibels to the entropy in bits per sample of the quantized sequence of coefficients.

In vector quantization, a library of vectors $\{y_k \in \mathbf{R}^{N'} : k = 1, \dots, K\}$ is used to approximate N'-coefficient segments of the signal or its transform. It is important that this

library be well adapted to the signal, and the classical implementations have constructed gigantic and growing libraries of vectors, with the resulting computational complexity problems.

The best-basis wavelet packet algorithm is similar to vector quantization, but has two significant differences. First, the library of vectors is fixed in advance, so there is no need to transmit it to the receiver, nor to enlarge it in an attempt to improve the fit to the signal. Secondly, there is a low complexity algorithm to compute correlations with all the elements in the library, and to select the best fit.

Compression by coefficient selection. The most basic compression method is that in which we discard negligibly small coefficients. We will express x as a linear combination of elements of our wavelet packet library, fix $\epsilon > 0$, and choose that basis subset which contains the fewest coefficients larger than ϵ. Then we store only those coefficients in this basis which are greater than ϵ.

This procedure is an orthogonal projection of x onto a lower dimensional subspace, but the choice of subspace could depend on x in general and the procedure is therefore not linear. Also, the description of the retained subspace will contain some of the information lost by the projection. Any choice of library will result in some improvement in the efficiency of the representation. We can adjust our library, however, to take advantage of *a priori* knowledge about the signal.

Compression by coalescing. A particular kind of vector quantization is replacing all the coefficients in a subspace of a decomposition with a single coefficient. For example, after a subband decomposition, we may replace the possibly complicated distribution of frequencies within each subband by a single representative, whose amplitude is chosen to conserve energy, and whose frequency is taken to be the mean frequency of the distribution. Or, we may code the distribution within the subband with its first few moments. Such methods are justified by models of perception such as [[]J], which assert that frequencies within subbands are partially masked by near neighbors.

12. Nonstandard Matrix Multiplication

Write $\mathcal{W}(\mathbf{R})$ for the collection of one-dimensional wavelet packets. Let w_{fsp} be a representative wavelet packet of frequency f at scale 2^{-s}, located at $p2^{-s}$. Quadrature mirror filters may be chosen so that $\mathcal{W}(\mathbf{R})$ is dense in many common function spaces. With

minimal hypotheses, $\mathcal{W}(\mathbf{R})$ will be dense in $L^2(\mathbf{R})$. Using the Haar filters $\{1/\sqrt{2}, 1/\sqrt{2}\}$ and $\{1/\sqrt{2}, -1/\sqrt{2}\}$, for example, produces $\mathcal{W}(\mathbf{R})$ which is dense in $L^p(\mathbf{R})$ for $1 < p < \infty$. Longer filters can generate smoother wave packets, so we can also produce dense subsets of Sobolev spaces, etc.

Basis subsets. Since $L^2(\mathbf{R}) \otimes L^2(\mathbf{R})$ is dense in $L^2(\mathbf{R}^2)$, the collection $\{w_X \otimes w_Y : w_X \in \mathcal{W}(X), w_Y \in \mathcal{W}(Y)\}$ is dense in the space of Hilbert-Schmidt operators. Call $\sigma \subset \mathbf{Z}^6$ a basis subset if the collection $\{w_{f_X s_X p_X} \otimes w_{f_Y s_Y p_Y} : (f_X, s_X, p_X, f_Y, s_Y, p_Y) \in \sigma\}$ forms a Hilbert basis.

Ordering wave packets. Wave packets w_{fsp} can be totally ordered. We say that $w < w'$ if $(m, n, k) < (m', n', k')$. The triplets are compared lexicographically, counting the scale parameter m as most significant.

Tensor products of wave packets inherit this total order. Write $w_X = w_{f_X s_X p_X}$, etc. Then we will say that $w_X \otimes w_Y < w'_X \otimes w'_Y$ if $w_X < w'_X$ or else if $w_X = w'_X$ but $w_Y < w'_Y$. This is equivalent to $(s_X, f_X, p_X, s_Y, f_Y, p_Y) < (m'_X, n'_X, k'_X, m'_Y, n'_Y, k'_Y)$ comparing lexicographically from left to right.

Define the *adjoint order* $<^*$ by exchanging X and Y indices, namely $w_X \otimes w_Y <^* w'_X \otimes w'_Y$ if and only if $w_Y \otimes w_X <^* w'_Y \otimes w'_X$. This is also a total order.

Projections. Let $\mathcal{W}^1$ denote the space of bounded sequences indexed by the three wave packet indices f, s, p. With the ordering above, we obtain a natural isomorphism between ℓ^∞ and $\mathcal{W}^1$. There is also a natural injection $J^1 : L^2(\mathbf{R}) \hookrightarrow \mathcal{W}^1$ given by $c_{fsp} = \langle v, w_{fsp} \rangle$, for $v \in L^2(\mathbf{R})$ and $w_{fsp} \in \mathcal{W}(\mathbf{R})$. If σ is a basis subset, then the composition J^1_σ of J^1 with projection onto the subsequences indexed by σ is also injective. J^1_σ is an isomorphism of $L^2(\mathbf{R})$ onto $l^2(\sigma)$, which is defined to be the square summable sequences of $\mathcal{W}^1$ whose indices belong to σ.

We also have a map $R^1 : \mathcal{W}^1 \to L^2(\mathbf{R})$ defined by $R^1 c(t) = \sum_{(f,s,p) \in \mathbf{Z}^3} c_{fsp} w_{fsp}(t)$. This map is defined and bounded on the closed subspace of $\mathcal{W}^1$ isomorphic to l^2 under the natural isomorphism mentioned above. In particular, R^1 is defined and bounded on the range of J^1_σ for every basis subset σ. The related restriction $R^1_\sigma : \mathcal{W}^1 \to L^2(\mathbf{R})$ defined by $R^1_\sigma c(t) = \sum_{(f,s,p) \in \sigma} c_{fsp} w_{fsp}(t)$ is a left inverse for J^1 and J^1_σ. In addition, $J^1 R^1_\sigma$ is a projection of $\mathcal{W}^1$. Likewise, if $\sum_i \alpha_i = 1$ and $R^1_{\sigma_i}$ is one of the above maps for each i, then $J^1 \sum_i \alpha_i R^1_{\sigma_i}$ is also a projection of $\mathcal{W}^1$. It is an orthogonal projection on any finite subset of $\mathcal{W}^1$.

Similarly, writing $\mathcal{W}^2$ for $\mathcal{W}^1 \times \mathcal{W}^1$, the ordering of tensor products gives a natural isomorphism between ℓ^∞ and $\mathcal{W}^2$. Objects in the space $L^2(\mathbf{R}^2)$, i.e., the Hilbert-Schmidt operators, inject into this sequence space $\mathcal{W}^2$ in the obvious way, namely $M \mapsto$

$\langle M, w_{f_X s_X p_X} \otimes w_{f_Y s_Y p_Y} \rangle$. Call this injection J^2. If σ is a basis subset of $\mathcal{W}^2$, then the composition J^2_σ of J^2 with projection onto subsequences indexed by σ is also injective. J^2_σ is an isomorphism of $L^2(\mathbf{R}^2)$ onto $\ell^2(\sigma)$, the square summable sequences of $\mathcal{W}^2$ whose indices belong to σ.

The map $R^2 : \mathcal{W}^2 \to L^2(\mathbf{R}^2)$ given by $R^2 c(x,y) = \sum c_{XY} w_X(x) w_Y(y)$, is bounded on that subset of $\mathcal{W}^2$ naturally isomorphic to ℓ^2. In particular, it is bounded on the range of J^2_σ for every basis subset σ.

We may also define the restrictions R^2_σ of R^2 to subsequences indexed by σ, defined by $R^2_\sigma c(x,y) = \sum_{(w_X, w_Y) \in \sigma} c_{XY} w_X(x) w_Y(y)$. There is one for each basis subset σ of $\mathcal{W}^2$. Then R^2_σ is a left inverse of J^2 and J^2_σ, and $J^2 R^2_\sigma$ is a projection of $\mathcal{W}^2$. As before, if $\sum_i \alpha_i = 1$ and σ_i is a basis subset for each i, then $J^2 \sum_i R^2_{\sigma_i}$ is also a projection of $\mathcal{W}^2$. It is an orthogonal projection on any finite subset of $\mathcal{W}^2$.

Applying operators to vectors. For definiteness, let X and Y be two named copies of $\mathbf{R}$. Let $v \in L^2(X)$ be a vector, whose coordinates with respect to wave packets form the sequence $J^1 v = \{\langle v, w_X \rangle : w_X \in \mathcal{W}(X)\}$.

Let $M : L^2(X) \to L^2(Y)$ be a Hilbert-Schmidt operator. Its matrix coefficients with respect to the complete set of tensor products of wave packets form the sequence $J^2 M = \{\langle M, w_X \otimes w_Y \rangle : w_X \in \mathcal{W}(X), w_Y \in \mathcal{W}(Y)\}$. We obtain the identity

$$\langle Mv, w_Y \rangle = \sum_{w_X \in \mathcal{W}(X)} \langle M, w_X \otimes w_Y \rangle \langle v, w_X \rangle \tag{12.1}$$

This identity generalizes to a linear action of $\mathcal{W}^2$ on $\mathcal{W}^1$ defined by

$$c(v)_{fsp} = \sum_{(f's'p')} c_{fspf's'p'} v_{f's'p'}. \tag{12.2}$$

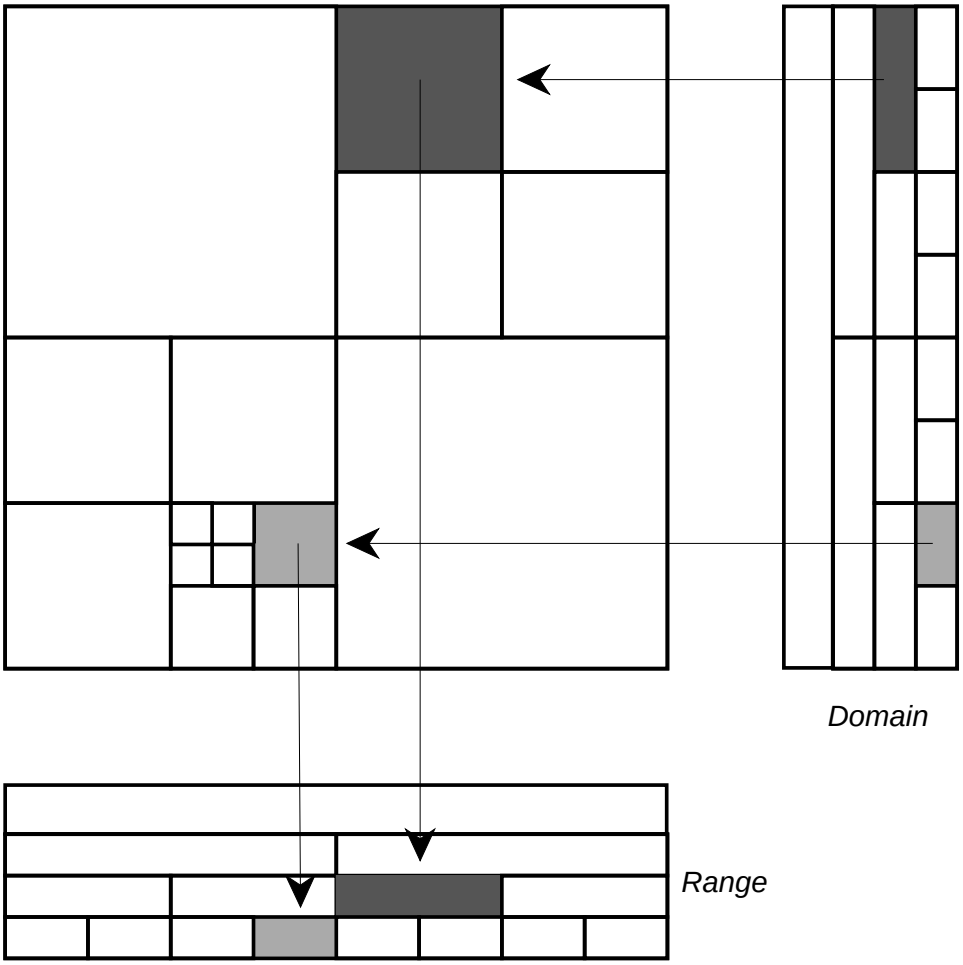
Now, images of operators form a proper submanifold of $\mathcal{W}^2$. Likewise, images of vectors form a submanifold $\mathcal{W}^1$. We can lift the action of M on v to these larger spaces via the commutative diagram

$$\begin{array}{ccc} \mathcal{W}^1 & \xrightarrow{J^2_\sigma M} & \mathcal{W}^1 \\ {\scriptstyle J^1}\uparrow & & \downarrow{\scriptstyle R^1} \\ L^2(\mathbf{R}) & \xrightarrow[M]{} & L^2(\mathbf{R}) \end{array} \tag{12.3}$$

The significance of this lift is that by a suitable choice of σ we can reduce the complexity of the map $J^2_\sigma M$, and therefore the complexity of the operator application.

Examples. We illustrate the above ideas by considerind nonstandard multiplication by matrices over $\mathbf{R}^8$. Using wavelet packets of isotropic dilations, we obtain the following diagram:

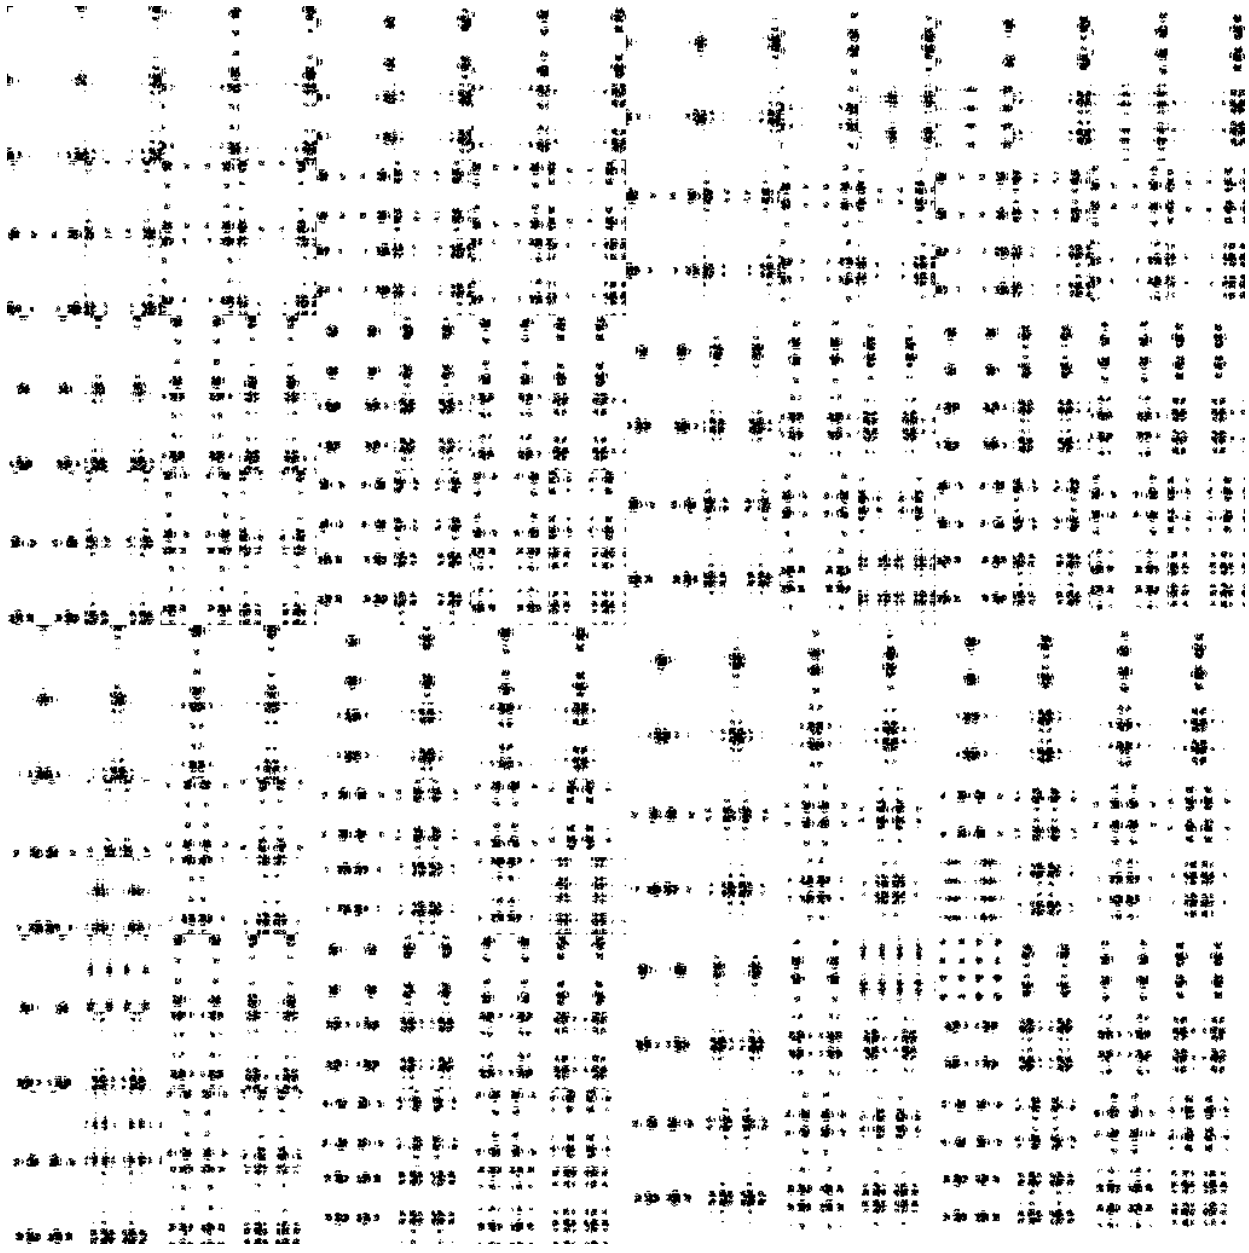
Nonstandard multiplication using restricted best bases for $\mathbf{R}^8 \times \mathbf{R}^8$



The domain vector is first injected into $\mathbf{R}^{8\log 8}$ by expansion into the complete wavelet packet tree at right. It then has its components multiplied by the best-basis coefficients in the matrix square at center, as indicated by the arrows. The products are summed into a range wavelet packet tree at bottom. The range tree is then projected onto the range space $\mathbf{R}^8$ by the adjoint of the wavelet packet expansion.

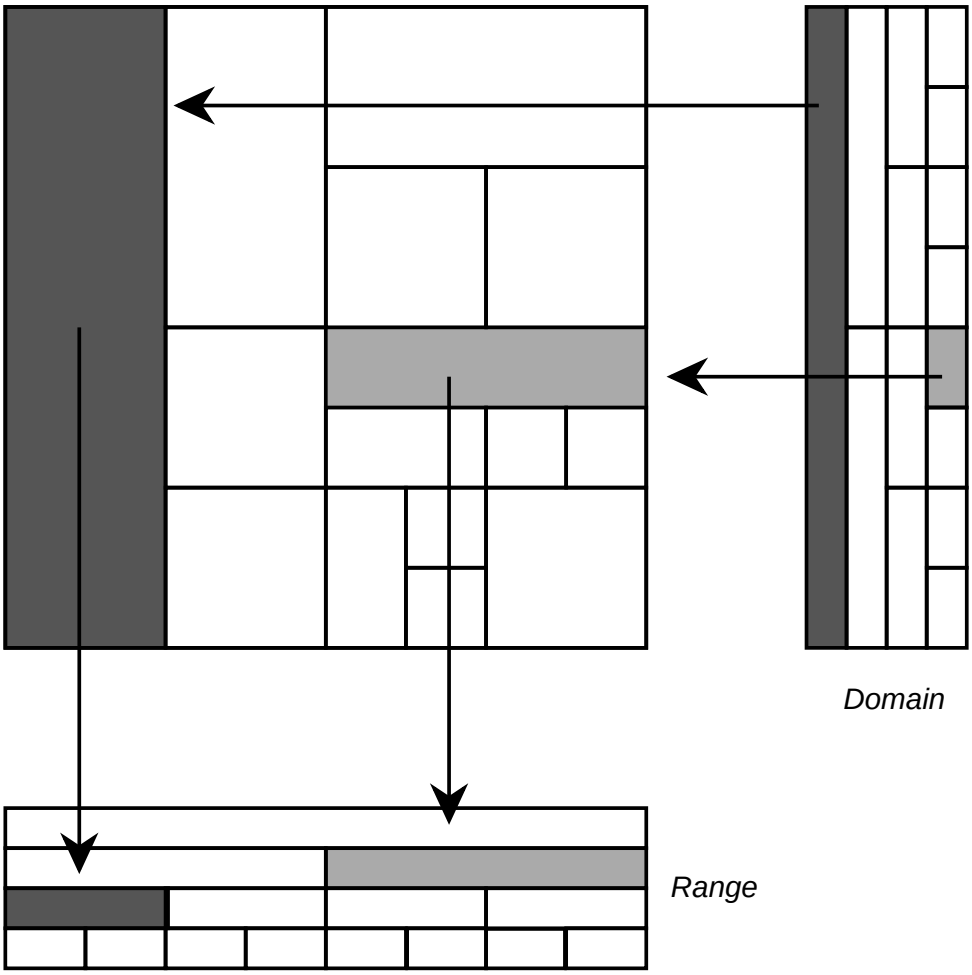
Below is the best-basis expansion for the matrix $m_{ij} = \sin \frac{\pi ij}{1024}$ where $0 \le i, j < 1024$. All coefficients smaller in magnitude than 1% of the largest have been set to zero; the survivors are depicted as black dots in the location they would occupy within the tree:

Restricted best basis for 1024 point sine transform, C30 filters



If we allow all tensor products of wavelet packets, finding the best basis for the matrix becomes harder, but the multiplication algorithm is practically the same:

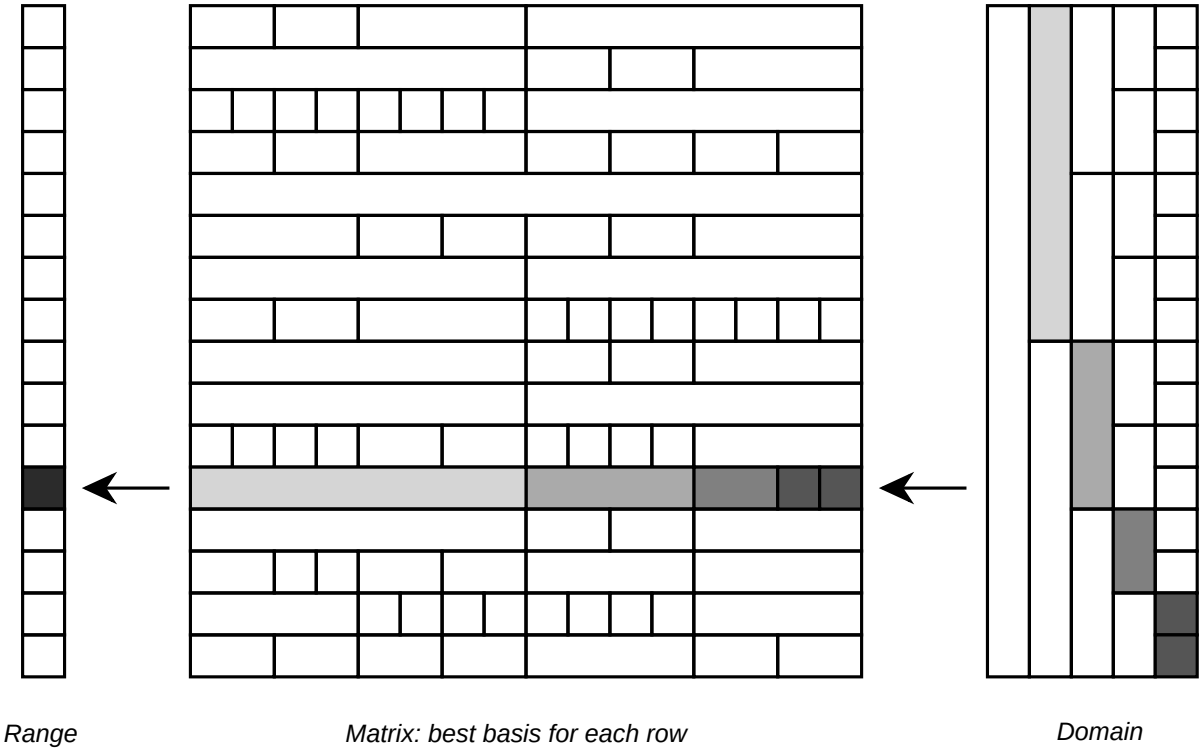
Nonstandard multiplication using best tensor product bases for $\mathbf{R}^8 \times \mathbf{R}^8$



As before, the domain vector is first injected into $\mathbf{R}^{8 \log 8}$, multiplied by the best tensor basis coefficients in the matrix square at center, as indicated by the arrows. The products are summed into the wavelet packet tree at bottom, and then projected onto $\mathbf{R}^8$.

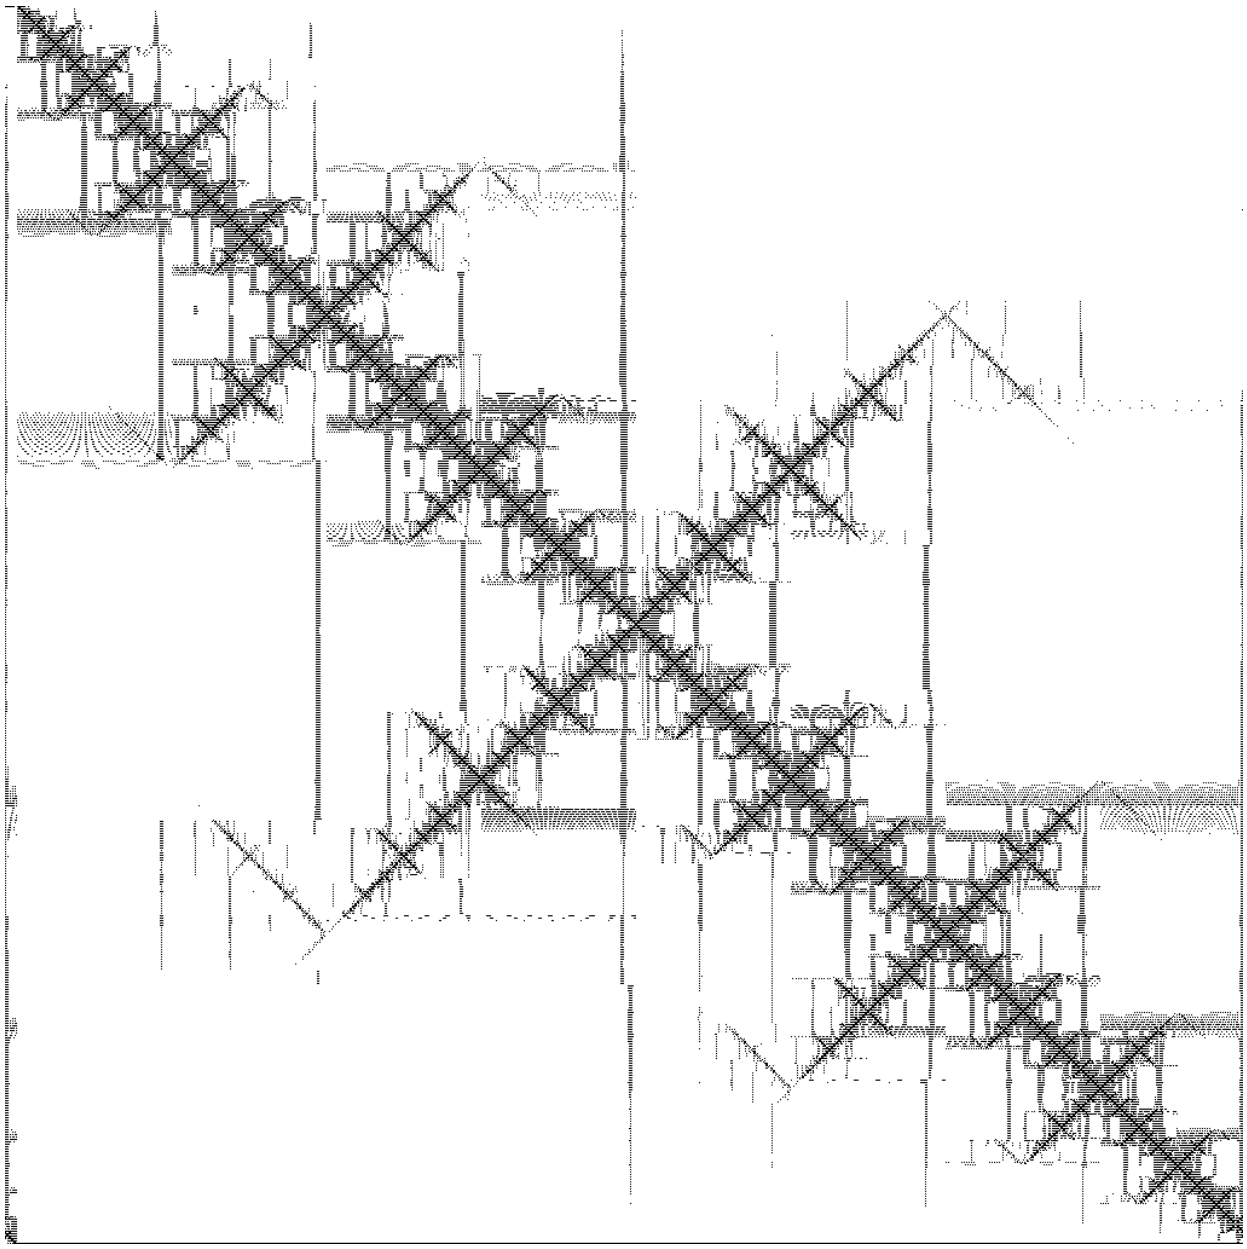
Rather than choose a basis set among the two-dimensional wavelet packets, we may also consider matrix multiplication to be a list of inner products. The complexity of the operation may be reduced by expressing each row of the matrix in its individual best basis, expanding the domain vector into the complete tree of wavelet packet coefficients, then evaluatin the inner products row by row, up to any desired accuracy. Schematically, this accelerated inner product may be depicted as below:

Nonstandard multiplication using best bases row by row for $\mathbf{R}^{16} \times \mathbf{R}^{16}$



A typical candidate for such a best-rows expansion is the 1024-point sine transform matrix, depicted below. In this schematic, enough coefficients were retained to preserve 99% of the Hilbert–Schmidt norm of each row, and the surviving coefficients are plotted as black dots in the locations they would occupy. Approximately 7% of the coefficients survived. No information is provided about the chosen bases in this picture:

Best row bases for 1024 point sine transform, C30 filters



Operation counts: transforming a vector. Suppose that M is a non-sparse operator of rank r. Ordinary multiplication of a vector by M takes at least $O(r^2)$ operations, with the minimum achievable only by representing M as a matrix with respect to the bases of its r-dimensional domain and range.

On the other hand, the injection J^2 will require $O(r^2[\log r]^2)$ operations, and each of J^1 and R^1 require $O(r \log r)$ operations. For a fixed basis subset σ of $\mathcal{W}^2$, the application

of $J^2_\sigma M$ to $J^1 v$ requires at most $\#|J^2_\sigma M|$ operations, where $\#|U|$ denotes the number of nonzero coefficients in U. We may choose our wavelet packet library so that $\#|J^2_\sigma M| = O(r^2)$. Thus the multiplication method described above costs an initial investment of $O(r^2[\log r]^2)$, plus at most an additional $O(r^2)$ per right-hand side. Thus the method has asymptotic complexity $O(r^2)$ per vector in its exact form, as expected for what is essentially multiplication by a matrix of order r.

We can obtain lower complexity if we take into account the finite accuracy of our calculation. Given a fixed matrix of coefficients C, write C_δ for the same matrix with all coefficients set to 0 whose absolute values are less than δ. By the continuity of the Hilbert-Schmidt norm, for every $\epsilon > 0$ there is a $\delta > 0$ such that $\|C - C_\delta\|_{HS} < \epsilon$. Given M and ϵ as well as a library of wavelet packets, we can choose a basis subset $\sigma \subset \mathcal{W}^2$ so as to minimize $\#|(J^2_\sigma M)_\delta|$. The choice algorithm has complexity $O(r^2[\log r]^2)$, as shown above. For a certain class of operators, there is a library of wavelet packets such that for every fixed $\delta > 0$ we have

$$\#|(J^2_\sigma M)_\delta| = O(r \log r), \tag{12.4}$$

with the constant depending, of course, on δ. Call this class with property (12.4) the *sparsifiable* Hilbert-Schmidt operators $\mathcal{S}$. By the estimate above, finite-precision multiplication by sparsifiable rank-r operators has asymptotic complexity $O(r \log r)$.

The best row basis algorithm is also asymptotically $O(r \log r)$ for sparsifiable matrices, since it requires $O(r \log r)$ operations to find the complete tree of wavelet packet coefficients for the domain vector, then $O(\log r)$ multiply-adds to evaluate each of the r inner products. Finding the best basis for each of the r rows of the matrix requires an initial investment of $O(r^2 \log r)$ operations.

Composing operators. Let X, Y, Z be three named copies of $\mathbf{R}$. Suppose that $M : L^2(X) \to L^2(Y)$ and $N : L^2(Y) \to L^2(Z)$ are Hilbert-Schmidt operators. We have the identity

$$\langle NM, w_X \otimes w_Z \rangle = \sum_{w_Y \in \mathcal{W}(Y)} \langle N, w_Y \otimes w_Z \rangle \langle M, w_X \otimes w_Y \rangle.$$

This generalizes to an action of $\mathcal{W}^2$ on $\mathcal{W}^2$, which is defined by the formula

$$c(d)_{fspf's'p'} = \sum_{f''s''p''} d_{fspf''s''p''} c_{f''s''p''f's'p'},$$

where c and d are sequences in $\mathcal{W}^2$. Using J^2, we can lift multiplication by N to an action

on these larger spaces via the commutative diagram

$$\begin{array}{ccc} \mathcal{W}^2 & \xrightarrow{J^2_\sigma N} & \mathcal{W}^2 \\ {\scriptstyle J^2}\uparrow & & \downarrow{\scriptstyle R^2} \\ L^2(\mathbf{R}^2) & \xrightarrow[N]{} & L^2(\mathbf{R}^2) \end{array}$$

Again, by a suitable choice of σ the complexity of the operation may be reduced to below that of ordinary operator composition.

Operation counts: composing two operators. Suppose that M and N are rank-r operators. Standard multiplication of N and M has complexity $O(r^3)$. The complexity of injecting N and M into $\mathcal{W}^2$ is $O(r^2[\log r]^2)$. The action of $J^2_\sigma N$ on $J^2 M$ has complexity $O(\sum_{fsp} \#|J^2_\sigma N_{YZ} : (f_Y, s_Y, p_Y) = (f, s, p)|\#|J^2 M_{XY} : (f_Y, s_Y, p_Y) = (f, s, p)|)$. The second factor is a constant $r \log r$, while the first when summed over all fsp is exactly $\#|J^2_\sigma N|$. Thus the complexity of the nonstandard multiplication algorithm, including the conjugation into the basis set σ, is $O(\#|J^2_\sigma N|\, r \log r)$. Since the first factor is r^2 in general, the complexity of the exact algorithm is $O(r^3 \log r)$ for generic matrices, reflecting the extra cost of conjugating into the basis set σ.

For the approximate algorithm, the complexity is $O(\#|(J^2_\sigma N)_\delta|\, r \log r)$. For the sparsifiable matrices, this can be reduced by a suitable choice of σ to a complexity of $O(r^2[\log r]^2)$ for the complete algorithm. Since choosing σ and evaluating J^2_σ each have this complexity, it is not possible to do any better by this method.

13. Fast Approximate Factor Analysis

Given an ensemble of random vectors, *principal factor analysis* seeks an ordered collection of uncorrelated basis vectors in which to most efficiently express the ensemble. Here, most efficiently means that the variance of the ensemble accumulates most rapidly as we take more and more of the basis vectors. The privileged basis most rapidly distinguishes the vectors in the ensemble, and is therefore useful for feature extraction and signal recognition. This basis is often called the Karhunen–Loéve basis.

Notation. Let $X = \{X_n : n = 1, \ldots, N\} \in \mathbf{R}^d$ be an ensemble of vectors. We suppose that d is very large and that X spans $\mathbf{R}^d$, implying $N \geq d$. The Karhunen–Loéve basis for

the ensemble consists of the eigenvectors of the symmetric positive definite autocovariance matrix $M = E(X \otimes X)$, or

$$M_{ij} = \frac{1}{N} \sum_{n=1}^{N} X_n(i) X_n(j).$$

It is known that coefficients with respect to the Karhunen–Loéve basis are independent random variables, and that they achieve the maximum linear transform coding gain or equivalently, the minimum entropy of any linear code used to transmit X.

Write $\overline{X} = E(X) = \frac{1}{N} \sum_{n=1}^{N} X_n$, and let $\sigma(X) \in \mathbf{R}^d$ be the vector of variances of the coefficients of X. Namely,

$$\sigma(X)(i) = \frac{1}{N} \sum_{n=1}^{N} [X_n(i) - \overline{X}(i)]^2.$$

We may assume without loss that $\overline{X} = 0$. Write $\mathrm{Var}(X)$ for the sum of the coefficients of $\sigma(X)$, which is the total variance of the ensemble X.

Let $U : \mathbf{R}^d \to \mathbf{R}^d$ be orthogonal and write $Y = UX$ for the map $Y_n = UX_n$. Since U is linear, $\overline{Y} = \overline{UX} = U\overline{X} = 0$, and since U is orthogonal, $\mathrm{Var}(X) = \mathrm{Var}(Y)$. Define the *transform coding gain* as in [JN] by the formula $G_{TC}(U) = \mathrm{Var}(UX) / \exp H(UX)$, where

$$H(X) = \frac{1}{d} \sum_{i=1}^{d} \log \sigma(X)(i).$$

$G_{TC}(UX)$ is maximized when $H(UX)$ is minimized, and H is the entropy of the direct sum of d independent Gaussian random variables with variances $\sigma(X)(i)$, $i = 1, \ldots, d$. The Karhunen–Loéve transform is a global minimum for H, and we will say that the best approximation to the Karhunen–Loéve transformation from a library $\mathcal{U}$ of orthogonal transformations is the minimum of $H(UX)$ with the constraint $U \in \mathcal{U}$.

The fast approximate Karhunen–Loéve algorithm. Notice that H is an information cost function as described above. We may create a large library of orthogonal bases by recursive quadrature mirror filter convolution-decimation, and use the best-basis search algorithm with H to find the best approximation to the Karhunen–Loéve basis. In the case $d = 8$ we have:

Complete wavelet packet expansion of X_i.

$X_i(1)$	$X_i(2)$	$X_i(3)$	$X_i(4)$	$X_i(5)$	$X_i(6)$	$X_i(7)$	$X_i(8)$
s_1	s_2	s_3	s_4	d_1	d_2	d_3	d_4
ss_1	ss_2	ds_1	ds_2	sd_1	sd_2	dd_1	dd_2
sss_1	dss_1	sds_1	dds_1	ssd_1	dsd_1	sdd_1	ddd_1

Then the sums of the squares are accumulated in an array of variances:

Variances of wavelet packet coordinates of X.

$\sum_i X_i^2(1)$	$\sum_i X_i^2(2)$	$\sum_i X_i^2(3)$	$\sum_i X_i^2(4)$	$\sum_i X_i^2(5)$	$\sum_i X_i^2(6)$	$\sum_i X_i^2(7)$	$\sum_i X_i^2(8)$
$\sum_i s_1^2$	$\sum_i s_2^2$	$\sum_i s_3^2$	$\sum_i s_4^2$	$\sum_i d_1^2$	$\sum_i d_2^2$	$\sum_i d_3^2$	$\sum_i d_4^2$
$\sum_i ss_1^2$	$\sum_i ss_2^2$	$\sum_i ds_1^2$	$\sum_i ds_2^2$	$\sum_i sd_1^2$	$\sum_i sd_2^2$	$\sum_i dd_1^2$	$\sum_i dd_2^2$
$\sum_i sss_1^2$	$\sum_i dss_1^2$	$\sum_i sds_1^2$	$\sum_i dds_1^2$	$\sum_i ssd_1^2$	$\sum_i dsd_1^2$	$\sum_i sdd_1^2$	$\sum_i ddd_1^2$

This tree of variances may be searched for the orthogonal basis which minimizes H. Let U be this basis, and write $\{U_i \in \mathbf{R}^d : i = 1, \ldots, d\}$ for the rows of U. We may suppose that these rows are numbered so that $\sigma(UX)$ is in decreasing order. Then we fix $\epsilon > 0$ and let d' be the smallest integer such that $\sum_{n=1}^{d'} \sigma(UX)(n) \geq (1-\epsilon)\mathrm{Var}(X)$. Then the projection of X onto the span of $U' = \{U_1 \ldots U_{d'}\}$ contains $1-\epsilon$ of the variance.

Statement of the main results. Call the orthogonal projection U' (associated to ϵ) the approximate Karhunen–Loéve transform with relative variance error ϵ. Already these d' vectors U' are a good basis for the ensemble X, but they may be further decorrelated by Karhunen–Loéve factor analysis. The algorithm is fast because we expect that even for small ϵ we will obtain $d' \ll d$. Counting operations:

(1) Expanding N vectors $X_n \in \mathbf{R}^d$ into wavelet packet coefficients: $O(Nd \log d)$.
(2) Summing squares into the variance tree: $O(d \log d)$.
(3) Searching the tree for a best basis: $O(d)$.
(4) Sorting the best basis vectors into decreasing order of importance: $O(d \log d)$.
(5) Transforming $U'X$ by Karhunen–Loéve: $O({d'}^3)$.

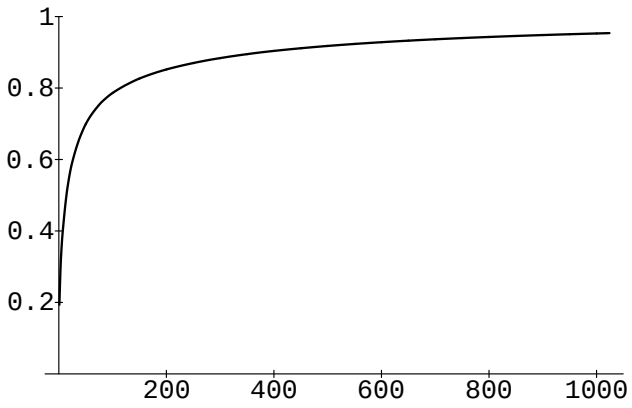
Indeed, the last step may not be necessary, since a large reduction in the number of parameters is already achieved by the orthogonal projection U'.

Application to the mug's gallery problem. Lawrence Sirovich provided 143 digitized 128x128x8bit pictures of Brown University students. These were already normalized with the pupils impaled on two fixed points near the center of the image. We first transformed the data to floating point numbers, computed average values for the pixels, and subtracted the average from each pixel to obtain "caricatures," or deviations from the average.

Each caricature was treated as a picture and expanded into 2 dimensional wavelet packets defined above, as described in [W]. The squares of the amplitudes were summed into a tree of variances, which was searched for a best-basis using the logarithm of energy information cost. Call this most-concentrated basis the *joint best basis* for the ensemble.

In the joint best basis, 400 coordinates (of 16384) contained more than 90% of the variance of the ensemble. Below we see the total variance on the first d' coordinates in the joint best basis, sorted in decreasing order, as a fraction of the total variance of the ensemble, for $1 \le d' \le 1024$.

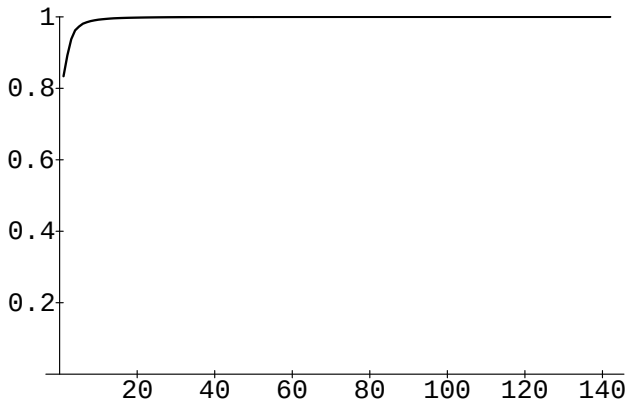
Accumulation of variance by joint best basis vectors



Using 1024 parameters captures more than 95% of the ensemble variance, but requires somewhat more computer power than is readily available on a desktop. A 400 parameter system can be analyzed on a common workstation in minutes. The top 400 coordinates were recomputed for each caricature and their autocovariance matrix over the ensemble was diagonalized by the LINPACK singular value decomposition routine.

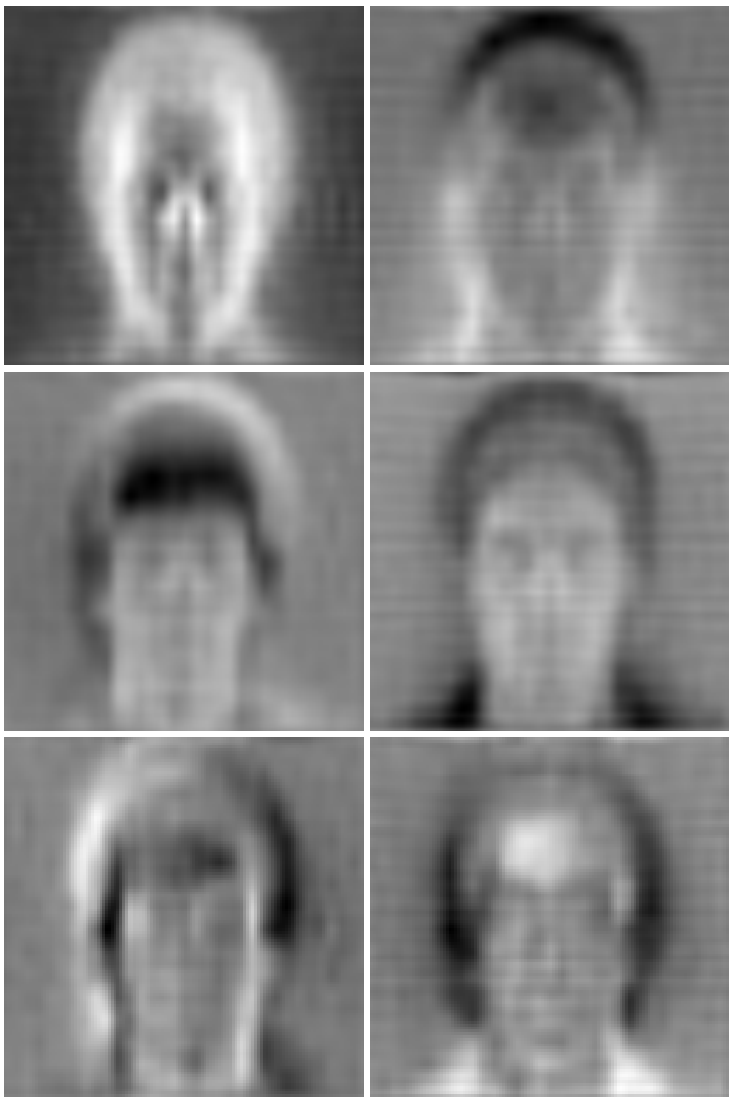
The graph below shows the total variance on the first d' coordinates in the Karhunen–Loéve basis, sorted in decreasing order, as a fraction of the total variance of the 400 joint best basis coefficients, for $1 \le d' \le 143$. The Karhunen–Loéve post-processing for this small ensemble concentrates 98% of the retained variance from the top 400 joint best-basis parameters into 10 coefficients.

Accumulation of variance by Karhunen–Loéve eigenvectors



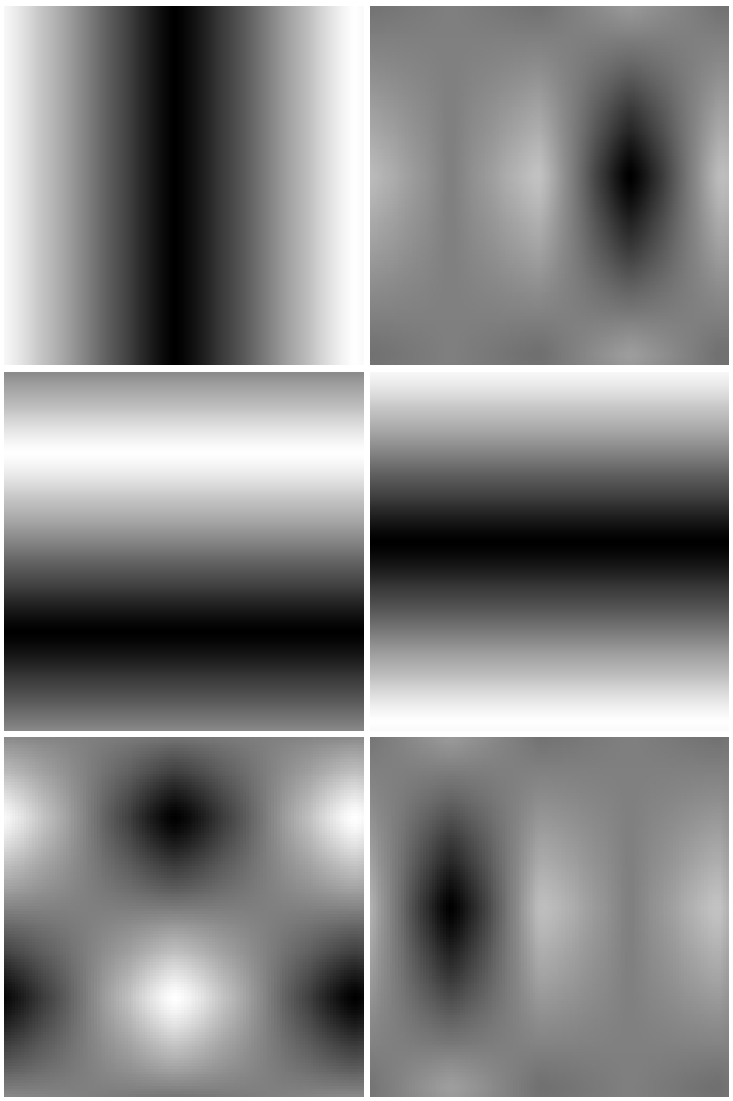
The top 6 "eigenfaces" with respect to the top 400 joint best-basis parameters look like ghostly caricatures. They were produced by reconstructing an image from each list of 400 joint best basis coefficients representing one of the 143 eigenvectors of the autocovariance matrix. They have been normalized to fill the dynamic range of the printing device.

Top 6 Karhunen–Loéve eigenfaces



The top 6 joint best basis wavelet packets are better from ananalytic point of view. One has a priori estimates on the size and derivatives, which is not available in the Karhunen–Loéve case. Whereas the Karhunen–Loéve basis functions look more or less like faces (or at least heads), the wavelet packets are abstract blobs which can be better localized at specific facial features:

Top 6 joint best basis vectors



Conclusions. Wavelet packet analysis reduces the number of parameters needed to perform approximate Karhunen–Loéve expansions. For a factor analysis "explaining" all but ϵ of the ensemble variance in a d parameter system, the complexity will be $O(d^2 \log d + {d'}^3)$, where $d' \ll d$. For accuracies of 1 or 2 significant digits, the analysis of systems of 16384 paramters can be performed on desktop computers in minutes.

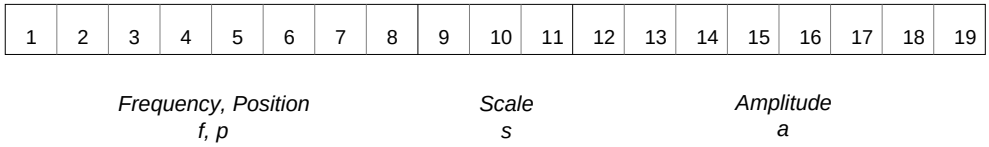
14. Speech Scrambling

Let $C = (a, f, s, p)$ denote the amplitude, frequency, scale, and position parameters of a wavelet packet coefficient, respectively. A sequence $\{C_i : i = 1, 2, \dots\}$ of such coefficients represents an encoded speech signal. The sequence has an associated distribution of values. One may transform $C_i \to C_i'$ so as to preserve this distribution. Signals reconstructed from C' rather than from C will be garbled, but in a way that makes the garbled signal hard to distinguish automatically from ordinary speech. In effect, the scrambled speech will seem like some other speaker using a different language.

The map $\pi : \{C_i\} \to \{C_i'\}$ may be taken to be invertible (or even a permutation of the indices n, m, k). The map can be determined by a public key encryption method. Two individuals may each have a key, so that they may share a conversation in confidence, with eavesdroppers facing a difficult problem unscrambling either half.

Objects to be scrambled. Suppose we digitally sample speech at 8012 8-bit samples per second, as in telephony. We divide this into 256-sample (31 ms) windows and find the best-basis representation, quantizing and retaining the 13 most energetic coefficients (around 5%) per window. These may be packaged as follows:

Tagged format for a 1-dimensional wavelet packet coefficient



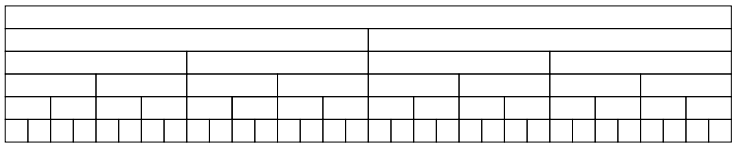
The result is $13 \times 19 = 247$ bits per window, or a rate of just under 8kbps without any entropy coding or removal of silences.

Feature-preserving permutations. Let C_i, $i = 1, 2, \dots$ represent a sequence of wavelet packet coefficients, which we shall take to represent a compressed speech signal. Although the transmitted coefficients will in general be further encoded, we will assume that this is done in a lossless manner so that the parameters a, s, f, p may be extracted.

To the σ-algebra $\mathcal{C}$ of ensembles $C = (a, f, s, p)$ we associate a probability measure $P : \mathcal{C} \to [0, 1]$. For definiteness, take $\mathcal{C}$ to be all triplets $(\cdot, f, s, p)$, where $0 \le s \le L$, $0 \le f < 2^m$, and $0 \le p < 2^{L-s}$ for some fixed maximum level L. We will treat the amplitude as an energy density and say that the probability of a coefficient is the average energy per unit time found at its index. This probability may be determined empirically

from a sufficiently long sample of speech. Then the probability density may be depicted as a graph over the tree of wavelet packet coefficients:

Schematic distribution of wavelet packet coefficients



Wavelet Packet Coefficients



Example Coefficient Probability Density

The above is a freehand schematic of the situation, and is much less complicated than is to be expected in practice. To create such a picture with actual data, we will accumulate $|a|^2$ in the coefficient position (f, s, p) for a reasonably long segment of real speech. Here is the author reciting part of "The Walrus and the Carpenter" by Lewis Carrol, compressed to 8 kbps by discarding small coefficients of the best-basis on windows of 512 samples in 62 ms:

Actual distribution of wavelet packet coefficients



We quantize the probability density into a small number of level regions with approximately equal numbers of coefficients in each region. Suppose that we choose 5 regions, as shown in the example. Encryption will consist of applying 5 permutations, one in each of the 5 regions.

The deviation between the statistical distribution of the original and encrypted coefficients may be made smaller by increasing the number of level regions. Of course, this

number must remain small compared to the number of wavelet packet coefficients, otherwise there will be only trivial permutations within the regions.

15. Practical Considerations

Certain practical issues arise when algorithms are implemented which have little to do with the mathematics of the transformation. We shall consider the spreading of compactly supported wavelets in the non-periodic case, and the memory requirements of the algorithm.

Aperiodic wavelets and wavelet packets. Suppose that $\{x_n : n \in \mathbf{N}\}$ is a finite sequence of length N in ℓ^2, i.e., all elements vanish except for $x_1, \ldots, x_N$. Then each pair of filter convolution-decimations produces two new finite sequences in ℓ^2, thereby imposing no boundary conditions on the transform. The convolution-decimation may then be applied recursively, and the resulting decomposition may be called the *aperiodic* wavelet-packet expansion. It produces a library of orthonormal bases, and is useful in practice because all the subspaces produced are finite dimensional.

Each of the images of the original sequence of length N has length $\lceil \frac{N}{2} \rceil + \lceil \frac{R}{2} \rceil - 1$, where R is the number of coefficients in the impulse response of the filter, and $\lceil x \rceil$ denotes the smallest integer greater than or equal to x. This count comes from the observation that there are $\lceil \frac{N}{2} \rceil$ adjacent pairs of elements with at least one being nonzero in the original sequence, and that there are $\lceil \frac{N}{2} \rceil + \lceil \frac{R}{2} \rceil - 1$ translates of the filter by 2 whose impulse response overlaps at least one of these pairs. We can express the number of coefficients resulting from L convolution-decimations as an iterated function:

Proposition 15.1. *Define $q(N) = \lceil \frac{N}{2} \rceil + \lceil \frac{R}{2} \rceil - 1$. Then after L iterations of convolution-decimation by filters of length R, the total number of coefficients from a sequence of length N is*

$$2^L q^L(N) = 2^L \overbrace{q(\ldots q(q(N))\ldots)}^{L \text{ iterations}} \quad \square$$

Corollary 15.2. *After L filter convolution-decimations, the total number of coefficients in the subband wavelet packet expansion at level L is no more than $N + (2^L - 1)R$.*

Proof. We can estimate $q(N) \leq \frac{N}{2} + \frac{R}{2}$. Iterating gives $q^L(N) \leq \frac{N}{2^L} + \frac{R}{2^L} + \frac{R}{2^{L-1}} + \cdots + \frac{R}{2}$. Multiplying by 2^L gives $N + (1 + \cdots + 2^{L-1})R = N + (2^L - 1)R$. $\square$

This may be sharpened in a typical special case:

Corollary 15.3. *If R is even, then after L filter convolution-decimations, the total number of coefficients in the subband wavelet packet expansion at level L is no more than $N + (2^L - 1)(R - 1)$.*

Proof. Under the hypothesis that R is even, we have the estimate that $q(N) \le \frac{N}{2} + \frac{R}{2} - \frac{1}{2}$. A calculation similar to the previous corollary now yields the sharper result. □

These results guarantee that the number of coefficients produced by the aperiodic wavelet packet algorithm does not grow too much:

Corollary 15.4. *If R is even and $2^L \le N$, then the aperiodic wavelet packet expansion at level L with filters of length R produces fewer than RN coefficients. Furthermore, the complete set of wavelet packet coefficients from the original signal down to level L includes no more than $(L+1)N + (R-1)(2^{L+1} - L - 2) \le (L+1)N + (R-1)(2N - L - 2)$ nonzero coefficients. For fixed R, this is $O(N \log N)$ as $N \to \infty$.*

Proof. Both formulas follow from the previous corollary by counting. □

Since the inverse of the wavelet packet algorithm is the adjoint, similar formulas hold for the number of coefficients in a signal reconstructed from an aperiodic expansion. For example, we may reconstruct the signal from a single subspace with only finitely many nonzero coefficients:

Proposition 15.5. *Using filters of length R, the signal reconstructed from a single subspace at level $L > 0$ containing N contiguous nonzero coefficients has at most $2^L N + (2^L - 1)(R - 2)$ nonzero components.*

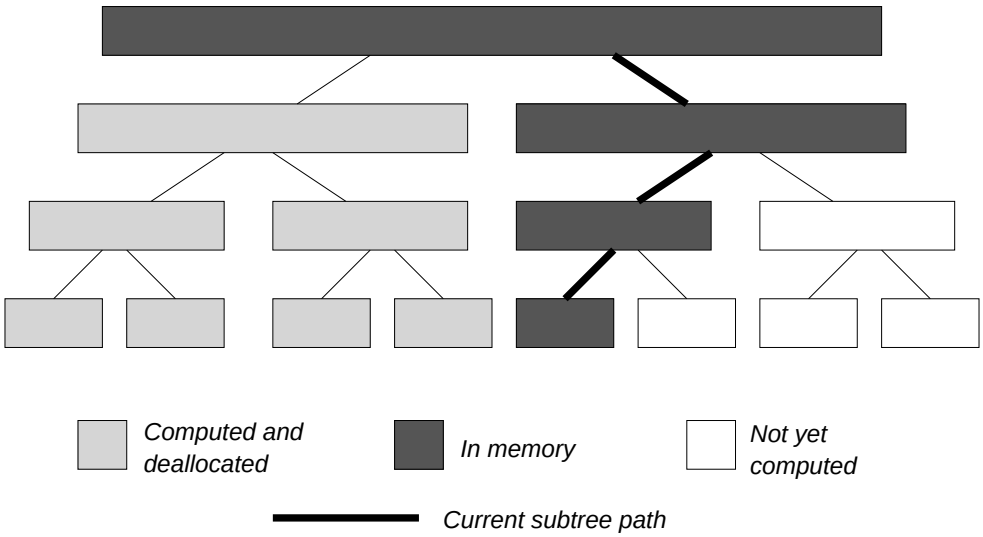
Proof. One application of the adjoint of convolution-decimation yields $2N + R - 2$ coefficients. Iterating L times gives $2^L N + 2^{L-1}(R-2) + \cdots + (R-2) = 2^L N + (2^L - 1)(R-2)$. □

With maximal overlap, the reconstruction from N coefficients in each of the 2^L subspaces at level L will also contain at most $2^L N + (2^L - 1)(R - 2)$ coefficients.

Trading space for time. Memory requirements for d-dimensional signals grow rapidly with d. If the signal can be decimated by 2 down to level L, then there must be at least 2^{Ld} samples all together. If d is large so that $N = 2^{Ld}$ is enormous even for moderate L, it may be impractical to store all LN wavelet packet coefficients. This problem can be overcome by trading off space for time, namely discarding the computed coefficients as soon as their information cost has been recorded. This requires that the coefficients be calculated recursively in depth-first pre-order, and that the memory used by a parent

subspace be freed as soon as the information cost of that parent has been compared to the cost of the best-basis among its descendents. The diagram below depicts an intermediate stage in such an algorithm in the 1-dimensional case.

Space-saving best basis search



Using this diagram one can easily prove that it takes no more than $\frac{2^d}{2^d-1}N$ memory locations (plus a constant overhead depending on d) to find a best basis for an N-sample signal of dimension d. The tradeoff is that each wavelet packet coefficient may have to be computed twice, at worst doubling the computation time.

The reconstruction algorithm is naturally organized to use no more than $2N$ memory locations for a signal of N samples.

Appendix 1: WPLab

The software program "WPLab," which produced the pictures of phase plane analyses, runs on NeXT computers and may be obtained by anonymous FTP (InterNet file transfer protocol) from `ceres.math.yale.edu [130.132.23.22]`, which is a site located within Yale University in New Haven, Connecticut, USA. Alternatively, please write or send electronic mail to the author (`victor@math.yale.edu`) to obtain a diskette.

Below is the "Help" panel from the program:

USING THE WAVELET PACKET LABORATORY

Reading a Signal From a File

WPLab can read text files which have the extension "`.asc`" appended to their names, and which contain just ASCII floating-point numbers.

Selecting the "Open" item from the application's main menu brings up a browser panel which allows the user to select any single file with the proper extension. The entire file is read into a double-precision floating-point array in memory, and the array is padded with zeroes up to an integer multiple of the longest available segment length.

The number of samples in the signal is displayed in the "Signal Length" text field. The signal file name becomes the title for the main window, and also for the miniwindow upon miniaturization.

Displaying a Signal Segment

Segments of the signal file are plotted in the rectangular view near the bottom of the window.

WPLab can display segments whose lengths are a power of 2, and starting at arbitrary offsets. Use the "Segment Length" radio buttons to select this length, and any combination of the offset form, slider, and buttons to set the index of the first displayed sample.

The buttons "Prev", "Next", "++", and "−−" adjust the offset. Their actions, respectively, are to subtract a segment length, add a segment length, add 1, and subtract 1. The program does the best it can given the signal length.

If a newly selected segment length is too long for the current offset, the offset is decreased to accommodate it.

Choosing a Quadrature Mirror Filter

There are 17 quadrature mirror filters (QMFs) available for wavelet packet analysis. They are identified by a letter ("B", "C", "D", or "V"), followed by (finite) impulse response length. For example, the Haar filters $\{\sqrt{\frac{1}{2}}, \sqrt{\frac{1}{2}}\}, \{\sqrt{\frac{1}{2}}, -\sqrt{\frac{1}{2}}\}$ are designated D2.

Preview a filter by clicking on its radio button. In the small "Wavelet" window will appear a plot of the mother wavelet associated to that QMF. This action enables the "Set QMF" and "Revert" buttons.

Roughly speaking, longer filters produce smoother wavelets and wavelet packets with better frequency resolution.

Click on the "Set QMF" button to register your choice and update the Phase Plane. This disables the "Set QMF" and "Revert" buttons.

Click on the "Revert" button to cancel previewing filters and return to the last registered QMFs. This will disable the "Set QMF" and "Revert" buttons.

The Phase Plane Representation of a Signal

The large square view contains that portion of the phase plane affected by the plotted segment. WPLab draws a rectangle in the phase plane for every modulated waveform in the basis chosen to represent the signal.

Each modulated waveform can be assigned 4 attributes: amplitude a, timescale s, frequency f, and position p. In a musical note, these correspond to loudness, duration, pitch, and the instant it is played.

Suppose that the signal segment has length $N = 2^n$. Coefficient (a, s, f, p) is displayed as the rectangle $[p2^s, (p+1)2^s] \times [f2^{n-s}, (f+1)2^{n-s}]$, shaded in proportion to a^2.

Because of the Heisenberg uncertainty principle, position and frequency cannot both be specified to arbitrary precision. The uncertainty of the frequency is 2^s, and the uncertainty in position is $2^(n-s)$. Thus each rectangle or "Heisenberg box" has area $2^n = N$. Since the total area of the displayed section of the phase plane is N^2, there are exactly N Heisenberg boxes in a disjoint cover of the section.

A library consists of all possible Heisenberg boxes, and bases from the library consist of certain disjoint covers of the phase plane by such rectangles.

Choosing a Basis

Wavelet Basis: this forces a display of the wavelet basis constructed with the given mother wavelet.

Best Level: this forces all of the Heisenberg boxes to have the same time scale. In particular, they must be congruent. There are $\log_2 N$ such bases for a segment of length N, and the one displayed has minimum entropy.

Best Basis: this minimizes entropy over all bases corresponding to disjoint dyadic covers of the segment. There are more than 2^N such bases for a segment of length N.

Printing the Window

The entire contents of the key window may be printed at full scale with the "Print" menu item.

References

[Al] Bradley Alpert, Dissertation, Yale University, 1990.

[Au] Pascal Auscher, *Symmetry properties for Wilson bases and new examples with compact support*, preprint, Department of Mathematics, Washington University, (St. Louis, Missouri) 1990.

[CM1] R. R. Coifman and Y. Meyer, *Nouvelles bases othonormées de $L^2(\mathbf{R})$ ayant la structure du système de Walsh*, preprint, Yale University, New Haven (1989).

[CM2] R. R. Coifman and Y. Meyer, *Remarques sur l'analyse de Fourier à fenêtre*, série I, C. R. Acad. Sci. Paris **312** (1991), 259–261.

[CMW] R. R. Coifman, Y. Meyer, and M. V. Wickerhauser, *Size properties of wavelet packets*, preprint, CEREMADE, Université Paris-Dauphine (1990) (to appear).

[CW] Ronald R. Coifman and M. Victor Wickerhauser, *Best-adapted wave packet bases*, preprint, Yale University (1990).

[D] Ingrid Daubechies, *Orthonormal bases of compactly supported wavelets*, Communications on Pure and Applied Mathematics **XLI** (1988), 909–996.

[DJJ] I. Daubechies, S. Jaffard, and J.-L. Journé, *A simple Wilson orthonormal basis with exponential decay*, SIAM J. Mathematical Analysis (to appear).

[DGM] I. Daubechies, A. Grossmann and Y. Meyer, *Painless nonorthogonal expansions*, J. Math. Phys. **27** (1986), 1271–1283.

[JN] Nuggehally S. Jayant and Peter Noll, *Digital coding of waveforms : principles and applications to speech and video*, Prentice-Hall, Englewood Cliffs, New Jersey, 1984.

[L] E. Laeng, *Une base orthonormale de $L^2(\mathbf{R})$, dont les éléments sont bien localisés dans l'espace de phase et leurs supports adaptés à toute partition symétrique de l'espace des fréquences*, série I, C. R. Acad. Sci. Paris **311** (1990), 677–680.

[Ma] Stephane G. Mallat, *A Theory for Multiresolution Signal Decomposition: The Wavelet Decomposition*, IEEE Transactions on Pattern Analysis and Machine Intelligence **11** (1989), 674–693.

[Ml] H. Malvar, *Lapped transforms for efficient transform/subband coding*, IEEE Trans. Acoustics, Speech, and Signal Processing **38** (1990), 969–978.

[Me] Yves Meyer, *Ondelettes et Opérateurs I,II*, Hermann, Paris, 1990.

[S] J. Strömberg, *A modified Haar system and higher order spline systems on $\mathbf{R}^n$ as unconditional bases for Hardy spaces*, *Conference inharmonic analysis in honor of Antoni Zygmund II*, W. Beckner et al. (Eds.), Wadsworth, Belmont, California, 1981, pp. 475–493.

[RY] K. R. Rao and P. Yip, *Discrete Cosine Transform*, Academic Press, New York, 1990.

[SW] Claude E. Shannon and Warren Weaver, *The Mathematical Theory of Communication*, The University of Illinois Press, Urbana, 1964.

[W] M. Victor Wickerhauser, *Picture compression by best-basis sub-band coding*, preprint, Yale University (New Haven, Connecticut) 1990.