

Optimality in the Design of Overcomplete Decompositions

Nick Kingsbury^a and H. Joel Trussell^b

^aSignal Processing & Comms. Group, Dept. of Engineering, Univ. of Cambridge, UK.

^bDept. Electrical & Computer Engineering, North Carolina State Univ., Raleigh, NC.

ABSTRACT

We lay a philosophical framework for the design of overcomplete multidimensional signal decompositions based on the union of two or more orthonormal bases. By combining orthonormal bases in this way, tight (energy preserving) frames are automatically produced. The advantage of an overcomplete (tight) frame over a single orthonormal decomposition is that a signal is likely to have a more sparse representation among the overcomplete set than by using any single orthonormal basis. We discuss the question of the relationship between pairs of bases and the various criteria that can be used to measure the goodness of a particular pair of bases. A particular case considered is the dual-tree Hilbert-pair of wavelet bases. Several definitions of optimality are presented along with conjectures about the subjective characteristics of the ensembles where the optimality applies. We also consider relationships between sparseness and approximate representations.

Keywords: Overcomplete decompositions, tight frames, unions of bases, dual-tree transforms, Hilbert pairs of bases, sparsity.

1. INTRODUCTION

The concept of overcomplete representations for signals has been around for some time. Frames, a particular class of overcomplete representations, have been defined since the 1950s. A very good introduction to this topic with many references to earlier work is the pair of tutorial papers by Kovačević and Chebira,^{1,2} so we do not repeat this historical material here. However there is still rather a lack of results in the area of optimal design of overcomplete decompositions of signals, and especially of work which takes account of a number of recent results in the area of sparse signal representations. Part of the problem has been the strong emphasis on non-redundant and orthogonal decompositions in the 1980s and 1990s (e.g. the discrete cosine transform and the discrete wavelet transform), because of their beautiful mathematical properties and highly efficient implementations. However more recently the potential advantages of overcomplete (redundant) decompositions have been recognized and/or rediscovered, so the question of optimal design of such systems is highly relevant.

This paper is somewhat tutorial in nature, but it also relates current general ideas about overcompleteness and sparsity to more specific cases involving Hilbert pairs of orthonormal bases (ONBs). It is in this latter area that we feel we have some new thoughts and results to offer.

2. TIGHT FRAMES AND UNIONS OF ORTHONORMAL BASES

Here we define some key mathematical concepts and set up notation for dealing with frames that are unions of orthonormal bases (ONBs). This rather brief analysis follows closely the fundamental ideas about frames, as presented recently by Kovačević and Chebira.¹ We define the basis vectors of the frame to be columns of the $n \times m$ matrix Φ such that any n -dimensional signal \mathbf{x} is represented using an m -dimensional coefficient vector \mathbf{y} by

$$\mathbf{x} = \Phi \mathbf{y} \tag{1}$$

It is simplest (but often not optimal – see later) to calculate \mathbf{y} from \mathbf{x} using an $m \times n$ matrix Φ^* , so that

$$\mathbf{y} = \Phi^* \mathbf{x} \quad \text{where } \Phi \Phi^* = I_n \tag{2}$$

Further author information: (Send correspondence to N.K.)

N.K.: E-mail: ngk@eng.cam.ac.uk, Telephone: +44 (0)1223 338514

H.J.T.: E-mail: hjt@eos.ncsu.edu, Telephone: 1 919 515 5126

We assume that $m \geq n$, and hence we cannot simply say that $\Phi^* = \Phi^{-1}$ unless $m = n$ and Φ is full rank. Note, in the non-redundant case when $m = n$, that $\Phi^* \Phi = I_m$; whereas in the redundant case when $m > n$, $\Phi^* \Phi$ becomes a projection matrix, projecting from any point in m -dimensional space to the nearest point within the n -dimensional range space of Φ^* .

In order for Φ and Φ^* to represent a frame, when \mathbf{y} is calculated from \mathbf{x} using equation (2) the following frame bound must be satisfied:

$$A \|\mathbf{x}\|^2 \leq \|\mathbf{y}\|^2 \leq B \|\mathbf{x}\|^2 \quad \text{for all } \mathbf{x}, \text{ and where } 0 < A \leq B \quad (3)$$

If $A = B$, the frame is known as a *tight frame* (TF). Furthermore, if $A = B = 1$, it is known as a *Parseval tight frame* (PTF) because $\|\mathbf{y}\|^2 = \|\mathbf{x}\|^2$ for any \mathbf{x} , and energy is preserved between the signal and transform domains, as in the standard Parseval theorem for Fourier Transforms.

Now

$$\|\mathbf{y}\|^2 = \mathbf{y}^H \mathbf{y} = \mathbf{x}^H \Phi^{*H} \Phi^* \mathbf{x} \quad (4)$$

Hence we see that if $\Phi \Phi^H = I_n$ so that $\Phi^* = \Phi^H$, then $\Phi^{*H} \Phi^* = \Phi \Phi^H = I_n$. Therefore

$$\|\mathbf{y}\|^2 = \mathbf{x}^H \Phi \Phi^H \mathbf{x} = \mathbf{x}^H \mathbf{x} = \|\mathbf{x}\|^2 \quad \text{for any } \mathbf{x} \quad (5)$$

and so we have a PTF. (Note that we use hermitian transpose, $[\cdot]^H$, for generality here, even though most matrices are likely to be purely real so transpose, $[\cdot]^T$, would be equivalent.)

If $m = n$ and $\Phi \Phi^H = I_n$, then this implies that Φ is an orthonormal matrix. Hence a non-redundant PTF must also be an ONB.

Now let us consider the case $m = Kn$, where Φ is a union of K n -dimensional ONBs, Φ_k , $k = 1 \dots K$, which are all different from each other:

$$\Phi = \frac{1}{\sqrt{K}} [\Phi_1 \quad \Phi_2 \quad \dots \quad \Phi_K] \quad (6)$$

For this case:

$$\Phi \Phi^H = \frac{1}{\sqrt{K}} [\Phi_1 \quad \dots \quad \Phi_K] \cdot \frac{1}{\sqrt{K}} \begin{bmatrix} \Phi_1^H \\ \vdots \\ \Phi_K^H \end{bmatrix} = \frac{1}{K} \sum_{k=1}^K \Phi_k \Phi_k^H = \frac{1}{K} K I_n = I_n \quad (7)$$

Hence a union of K ONBs, scaled by $1/\sqrt{K}$, produces a PTF. This is not unexpected as each ONB is energy preserving.

It is worth noting here that one or more additional stages of ONB conversions may follow the first stage of K parallel ONBs. Typically the ONBs of the new stage(s) operate on vectors which are samples from across two or more of the first-stage ONB outputs, thus cross-connecting outputs from the first stage. Energy is still preserved through these later ONBs so the overall Φ still represents a PTF.

3. KEY PROPERTIES OF A PARSEVAL TIGHT FRAME

How do we choose a good Φ or a good set of $\Phi_1 \dots \Phi_K$?

We can get some important hints on this from the ideas underlying principal components analysis (PCA), which tells us how to get an optimal Φ when $m = n$ and Φ comprises only one ONB. PCA works as follows.

Given p samples of \mathbf{x} , where typically $p \gg n$, we form a matrix

$$X = [\mathbf{x}_1 \quad \mathbf{x}_2 \quad \dots \quad \mathbf{x}_p] \quad (8)$$

and then calculate the symmetric covariance matrix (strictly an auto-correlation matrix)

$$S_X = \frac{1}{p} X X^H \quad (9)$$

The equivalent matrix to X in the transform domain, assuming Φ represents a PTF, is then

$$Y = [\mathbf{y}_1 \ \mathbf{y}_2 \ \cdots \ \mathbf{y}_p] = \Phi^H X \quad (10)$$

and so the covariance in the transform domain is

$$S_Y = \frac{1}{p} Y Y^H = \frac{1}{p} \Phi^H X X^H \Phi = \Phi^H S_X \Phi \quad (11)$$

Assuming $\Phi \Phi^H = I_n$ (as Φ represents a PTF),

$$\Phi S_Y \Phi^H = \Phi \Phi^H S_X \Phi \Phi^H = S_X \quad (12)$$

PCA theory says that an optimum transform is one which diagonalizes S_Y . An explanation for this, based on minimizing the off-diagonal energy in S_Y , is given in the appendix. Hence if $m = n$, Φ and S_Y contain the eigenvectors and eigenvalues respectively of S_X .

But what if $m > n$? Then the problem becomes more difficult.

For example, let $K = 2$ so $m = 2n$ and $\Phi = \frac{1}{\sqrt{2}} [\Phi_1 \ \Phi_2]$ where Φ_1 and Φ_2 each represent an ONB. Then

$$S_Y = \Phi^H S_X \Phi = \frac{1}{\sqrt{2}} \begin{bmatrix} \Phi_1^H \\ \Phi_2^H \end{bmatrix} S_X \frac{1}{\sqrt{2}} [\Phi_1 \ \Phi_2] = \frac{1}{2} \begin{bmatrix} \Phi_1^H S_X \Phi_1 & \Phi_1^H S_X \Phi_2 \\ \Phi_2^H S_X \Phi_1 & \Phi_2^H S_X \Phi_2 \end{bmatrix} \quad (13)$$

Hence S_Y is diagonalized if $\Phi_1^H S_X \Phi_1$ and $\Phi_2^H S_X \Phi_2$ are both diagonal and if $\Phi_1^H S_X \Phi_2 = (\Phi_2^H S_X \Phi_1)^H = O_n$, where O_n is a matrix of zeros. This is difficult for 2 reasons:

- It implies that Φ_1 and Φ_2 must *both* be eigenvector decompositions of S_X , which is not usually possible if they are to be significantly different from each other;
- It also implies that Φ_1^H times R_X must have zero correlation with Φ_2^H times R_X , where $R_X R_X^H = S_X$, (ie R_X is a square root of S_X).

But we can make progress if we allow the vectors in Y to be calculated some other way than using equation (2), while still ensuring that equation (1) is satisfied (for example by a sparsity inducing method such as Matching Pursuits – see next section). If this produces an approximately diagonal S_Y with rapidly decaying elements, such that $\Phi S_Y \Phi^H \simeq S_X$, then this implies that we can get a good approximation to X that is sparse in the \mathbf{y} domain, by eliminating the components of Y that correspond to low-valued elements of S_Y .

Hence we should choose Φ such that this can be achieved!

Let us consider the case of other values for \mathbf{y} which still result in the same \mathbf{x} when projected using $\Phi \mathbf{y}$. They can be expressed as

$$\mathbf{y} = \mathbf{y}_r + \mathbf{y}_n \quad \text{where } \mathbf{y}_r = \Phi^H \mathbf{x} \text{ and } \Phi \mathbf{y}_n = \mathbf{0} \quad (14)$$

Hence

$$\Phi \mathbf{y} = \Phi \mathbf{y}_r + \Phi \mathbf{y}_n = \Phi \Phi^H \mathbf{x} + \mathbf{0} = \mathbf{x}, \text{ assuming that } \Phi \Phi^H = I_n \text{ (i.e. } \Phi \text{ is a PTF)}. \quad (15)$$

\mathbf{y}_r is known as the *range-space* component of \mathbf{y} , while \mathbf{y}_n is its *null-space* component. Note that \mathbf{y}_n is always orthogonal to \mathbf{y}_r , since, if it were not so, \mathbf{y}_n would affect the value of $\mathbf{x} = \Phi \mathbf{y}$, which it does not do by its definition above.

Let Θ be an $(m - n) \times m$ matrix which generates the null-space components of Φ , so that $\mathbf{y}_n = \Theta^H \mathbf{z}$ where \mathbf{z} may be any $(m - n)$ -dimensional vector. In order for \mathbf{y}_n to be always orthogonal to \mathbf{y}_r , all the rows of Θ must be orthogonal to all the rows of Φ . In other words

$$\text{since } \Phi \mathbf{y}_n = \mathbf{0}, \quad \Phi \Theta^H \mathbf{z} = \mathbf{0} \quad \forall \mathbf{z}, \quad \text{and so } \Phi \Theta^H = \mathbf{O} \quad (16)$$

where \mathbf{O} is an $n \times (m - n)$ matrix of zeros.

This demonstrates the orthogonality between the rows of Φ and Θ , or between the range-space and null-space components of \mathbf{y} . We note that in general \mathbf{y} has m degrees of freedom, while $\mathbf{y}_r = \Phi^H \mathbf{x}$ has only the n degrees of freedom of \mathbf{x} , and $\mathbf{y}_n = \Theta^H \mathbf{z}$ has the $(m - n)$ degrees of freedom of \mathbf{z} . Hence we may write

$$\mathbf{y} = \mathbf{y}_r + \mathbf{y}_n = \Phi^H \mathbf{x} + \Theta^H \mathbf{z} = \begin{bmatrix} \Phi \\ \Theta \end{bmatrix}^H \begin{bmatrix} \mathbf{x} \\ \mathbf{z} \end{bmatrix} = \Psi^H \begin{bmatrix} \mathbf{x} \\ \mathbf{z} \end{bmatrix} \quad (17)$$

where we have formed the $m \times m$ matrix

$$\Psi = \begin{bmatrix} \Phi \\ \Theta \end{bmatrix} \quad (18)$$

If Θ and Φ both represent PTFs such that $\Theta\Theta^H = I_{m-n}$ and $\Phi\Phi^H = I_n$, then Ψ has the property that

$$\Psi\Psi^H = \begin{bmatrix} \Phi \\ \Theta \end{bmatrix} \begin{bmatrix} \Phi^H & \Theta^H \end{bmatrix} = \begin{bmatrix} \Phi\Phi^H & \Phi\Theta^H \\ \Theta\Phi^H & \Theta\Theta^H \end{bmatrix} = \begin{bmatrix} I_n & O \\ O^H & I_{m-n} \end{bmatrix} = I_m \quad (19)$$

Hence Ψ represents an m -dimensional ONB, of which Φ and Θ are two parts, comprising n and $(m - n)$ basis functions respectively. This neatly illustrates Naimark's theorem,^{1,3} which tell us that 'every PTF can be realized as a projection of an ONB from a larger space'.

Returning now to the double ONB (DONB) case when $m = 2n$ and $\Phi = \frac{1}{\sqrt{2}} [\Phi_1 \ \Phi_2]$, then we can see that a valid choice for Θ is $\Theta = \frac{1}{\sqrt{2}} [\Phi_1 \ (-\Phi_2)]$, because

$$\Phi\Theta^H = \frac{1}{2} [\Phi_1 \ \Phi_2] \begin{bmatrix} \Phi_1^H \\ -\Phi_2^H \end{bmatrix} = \frac{1}{2} (\Phi_1\Phi_1^H - \Phi_2\Phi_2^H) = \frac{1}{2} (I_n - I_n) = O \quad (20)$$

Hence S_Y is now

$$\begin{aligned} S_Y &= S_{Y_r} + S_{Y_n} = \Phi^H S_X \Phi + \Theta^H S_Z \Theta \\ &= \frac{1}{2} \begin{bmatrix} \Phi_1^H S_X \Phi_1 & \Phi_1^H S_X \Phi_2 \\ \Phi_2^H S_X \Phi_1 & \Phi_2^H S_X \Phi_2 \end{bmatrix} + \frac{1}{2} \begin{bmatrix} \Phi_1^H S_Z \Phi_1 & -\Phi_1^H S_Z \Phi_2 \\ -\Phi_2^H S_Z \Phi_1 & \Phi_2^H S_Z \Phi_2 \end{bmatrix} \\ &= \frac{1}{2} \begin{bmatrix} \Phi_1^H (S_X + S_Z) \Phi_1 & \Phi_1^H (S_X - S_Z) \Phi_2 \\ \Phi_2^H (S_X - S_Z) \Phi_1 & \Phi_2^H (S_X + S_Z) \Phi_2 \end{bmatrix} \end{aligned} \quad (21)$$

If $S_Z = S_X$, the two off-diagonal terms become zero, and if Φ_1 and Φ_2 both (approximately) diagonalize $(S_X + S_Z)$, then S_Y will be diagonalized. Thus we can overcome the second problem encountered in equation (13).

We shall now consider how sparsity can be induced in \mathbf{y} and how this affects the choice of Φ .

4. SPARSITY

Sparsity has been proposed in recent years as a key aim for many good signal processing algorithms. It allows signals to be separated efficiently from noise (denoising), and also permits different signal components to be separated from each other (source separation). Furthermore it can act as an effective tool for regularization of ill-conditioned processes such as encountered in deconvolution and in compressed sensing.

PCA has long been recognized as a way of achieving a degree of sparsity by largely linear methods, where Φ is a non-redundant transform ($m = n$) selected so as to project \mathbf{x} in directions that maximise variance and minimise cross correlations between the components of \mathbf{y} (off-diagonal energy in S_Y).

When Φ is allowed to become redundant ($m > n$), each signal vector \mathbf{x} no longer has a unique transform-domain vector \mathbf{y} . The minimum energy solution for a PTF is the one in which $\mathbf{y} = \mathbf{y}_r = \Phi^H \mathbf{x}$ and thus lies in the range-space of Φ . Hence \mathbf{y} contains no null-space components \mathbf{y}_n . This minimizes energy because equation (26), in the appendix, shows that the total energy of \mathbf{y} is the sum of the range-space energy, which equals the energy of \mathbf{x} , and the null-space energy, which is zero when \mathbf{y}_n is zero. However this is not in general the sparsest

\mathbf{y} to give $\mathbf{x} = \Phi\mathbf{y}$. If Φ comprises a union of K ONBs, as in equation (6), then the minimum energy approach always builds the reconstructed signal from the sum of K equal-amplitude components, each one created by analysing and reconstructing \mathbf{x} with just one of the ONBs Φ_k . This means that every feature of the signal must be represented equally in \mathbf{y} by each of the K ONBs, shown in the summation terms in equation (7). Clearly this is not going to give a very sparse representation.

A much better approach is to represent each feature by just one of the ONBs, and to use the ONB whose bases match the feature most closely (i.e. with fewest terms). This does not require any change to the reconstruction rule, $\mathbf{x} = \Phi\mathbf{y}$, but it does require significant non-linearities to be introduced to the way that \mathbf{y} is calculated from \mathbf{x} . A well-known greedy algorithm that implements a near-optimal solution to this problem is orthogonal matching pursuits (OMP),⁴ in which basis vectors are found, one at a time, in decreasing order of their ability to reduce the error between the input signal and its sparse approximation using the basis vectors found so far. This is a rather inefficient algorithm for finding many components in large datasets.

A more efficient approach is iterative threshold shrinkage, in which many new components are updated at each iteration of the algorithm as they gradually appear above a threshold which is being slowly reduced (shrunk) as the iterations proceed. A good general overview of iterative thresholding methods can be found in Elad et al.⁵ A method, closely related to these, is stagewise OMP (StOMP) by Donoho et al.⁶

The overall effect of these sparsity-inducing algorithms is to create vectors \mathbf{y} that minimize a cost function $J(x)$ which is some measure of sparsity: typically $J(x)$ comprises the L_p -norm of \mathbf{y} where $0 \leq p \leq 1$, plus an error measure, typically $\|\mathbf{x} - \Phi\mathbf{y}\|^2$. This produces \mathbf{y} , whose range-space component \mathbf{y}_r is close to $\Phi^H\mathbf{x}$ such that $\Phi\mathbf{y}$ closely approximates \mathbf{x} , and whose null-space component \mathbf{y}_n is selected so that $\mathbf{y} = \mathbf{y}_r + \mathbf{y}_n$ is sparse. The \mathbf{y}_n component increases the energy of \mathbf{y} but usually in a relatively controlled way since the L_p -norm component of the cost function tends to penalize excessively large \mathbf{y} that also minimize $J(x)$. We shall now look at how big the energy of \mathbf{y}_n can get.

Let Φ be a union of K ONBs, and let the signal \mathbf{x} be very sparse and comprise just a single unit-energy component that is one of the $m = Kn$ basis vectors, say the i^{th} column $\phi_{k,i}$ from Φ_k . Then the sparsest \mathbf{y} will contain just one coefficient, corresponding to $\phi_{k,i}$ and with amplitude \sqrt{K} so as to overcome the $1/\sqrt{K}$ factor in equation (6). This \mathbf{y} will have an energy of K times the energy of \mathbf{x} , whereas the minimum energy solution for \mathbf{y} will have the same energy as \mathbf{x} because Φ is a PTF. Hence, in this case, the null-space component \mathbf{y}_n will have $K - 1$ times the energy of the range-space component \mathbf{y}_r (and of \mathbf{x}). Because this is the sparsest scenario for \mathbf{y} (just one non-zero component), this represents the worst-case energy amplification by the sparsity-inducing analysis process, which is therefore by a factor of K for a PTF that is a union of K ONBs.

For the DONB case considered earlier, $K = 2$, and so the maximum energy of \mathbf{y}_n equals the energy of \mathbf{y}_r , which equals the energy of \mathbf{x} . Typically in a region of the signal which is closely approximated by a single basis vector from Φ , then the energy of \mathbf{y}_n in this region will be close to this maximum; whereas in other regions where the signal corresponds to a number of correlated components superimposed, then the energy of \mathbf{y}_n will be lower than this.

It is worth noting that the situation discussed at the end of the previous section, in connection with equation (21) regarding the need for $S_Z \approx S_X$, would produce approximately equal energies for \mathbf{y}_n and \mathbf{y}_r , which are $\text{trace}(S_Z)$ and $\text{trace}(S_X)$ respectively (see the appendix). In fact the sparsity-inducing process should result in reduced correlations between the components of \mathbf{y} for Φ_1 and those for Φ_2 , and thus automatically give low energy in the regions of S_Y corresponding to the Φ_1 and Φ_2 interactions (the upper-right and lower-left quarters). The reason for this is that a good sparsity-inducing process tends to use either a component from Φ_1 or a similar component from Φ_2 to represent a given feature of the signal, but it prefers not to use both together. Hence it tends to give much lower correlation between these two components in S_Y than would be the case with an analysis process that just produced $\mathbf{y} = \mathbf{y}_r = \Phi\mathbf{x}$. Therefore a good sparsity-inducing process will tend to minimize the upper-right and lower left terms in equation (21).

For a more redundant union of ONBs, where $K > 2$, a good sparsity-inducing process will tend to minimize all the cross-correlation terms in S_Y that are of the form $\Phi_k^H S \Phi_l$ where $k \neq l$.

Now we consider how these ideas about sparsity may be used to help choose a good set of ONBs for Φ . We have seen that the desirability of having a near-diagonal covariance matrix S_Y in the transform domain remains

true even when the redundancy is quite large $m \gg n$, as long as sparsity-inducing analysis methods are employed to create the most useful transform domain representations.

One important class of overcomplete systems are those designed for signals which have several distinct types of component. For example many audio signals have segments which are strongly tonal, such as vowel sounds in speech or melodic sections of music, while other segments tend to be much more impulsive, such as consonant (plosive) sounds in speech or percussive sections of music. With such signals it makes great sense for the ONBs of an overcomplete system to be chosen such that one or more ONBs are optimized for one type of sound while other ONBs are optimized for another type of sound. Sparsity-inducing algorithms should then pick the basis functions appropriate for each type of sound and automatically adapt the decomposition to the changing statistics of the input signal.

We will not consider such systems further here, as the choice of ONBs depends heavily on the type of signal being analyzed, and in general, existing techniques for designing these ONBs either already exist (particularly if it is just one ONB per signal type), or the results presented later can be used separately for each signal type. So we will concentrate on the problem of designing overcomplete systems for signals with statistics that do not change substantially with time.

Considering the DONB case again, and referring back to equation (21), we shall assume that the sparsity-inducing analysis method produces $S_Z \approx S_X$, but that this is not an exact equality. Looking at the final matrix result of this equation, we can infer the following properties about the ONBs, Φ_1 and Φ_2 :

- $\Phi_1^H S_X \Phi_1$ should be approximately diagonal.
- $\Phi_2^H S_X \Phi_2$ should also be approximately diagonal.
- $\Phi_1^H (S_X - S_Z) \Phi_2$ should be approximately zero.

The first two properties require both Φ_1 and Φ_2 to be close to eigenvector solutions of S_X (and of S_Z). The third property shows that it is desirable for S_Z to approximately equal S_X , but that it will be helpful if Φ_1 and Φ_2 are as uncorrelated with each other as possible.

A useful way of framing the problem of choosing Φ is to ask the question: Assuming diagonal S_Y is desirable, can we obtain a good approximation to S_X for typical signals from it? This leads to the importance of shift invariant statistics which apply to many typical signals.

5. SHIFT INVARIANCE

A key feature of many data sources is that their first and second order statistics are *shift invariant*. In practice this means that a given type of waveform feature can occur with equal probability at any point along the data stream (i.e. at any time, if the data is time dependent). It also means that any shift of the complete data stream is an equally valid data stream, ignoring end effects.

Within the framework we have been using in earlier sections, shift-invariance of \mathbf{x} means that the matrix of training data X can contain all shifts of the input data stream. Hence the i^{th} column of X will be $\mathbf{x}_i = [x_i \ x_{i+1} \ \dots \ x_{i+n-1}]^T$ where the x_j are samples from the input data. Shift invariance then means that the mean and the covariance of X will be independent of shift, so the auto-correlation matrix S_X will be invariant along its diagonals (i.e. of symmetric Toeplitz form). Hence

$$(S_X)_{i,j} = (S_X)_{i+c,j+c} \quad \forall \ i, j, i+c, j+c \text{ within the bounds of } S_X \quad (22)$$

We define S_X as being shift invariant if this applies.

Note that there may be some departure from this diagonal invariance in the vicinity of the boundaries of S_X but we shall ignore this here. We shall assume that S_X is large enough that its elements decay to zero along most of the boundaries except near the upper-left and lower-right corners.

If S_X is shift invariant, what does this imply about optimal choices for Φ and the resulting matrix S_Y ?




Figure 1. (a) Diagonal matrix S_Y for a 4-level discrete wavelet transform (DWT), assuming a $1/\omega^2$ spectral decay of energy and $m = n = 128$; (b) matrix S_X corresponding to this S_Y when Φ represents a DWT; (c) the equivalent S_Y for a dual-tree complex wavelet transform (DTCWT) where $m = 2n = 256$; (d) S_X when Φ represents a DTCWT.

Since we are aiming to find an optimal PTF Φ which diagonalizes S_Y , we are looking for equation (12) to be satisfied, i.e. $\Phi S_Y \Phi^H = S_X$ where S_Y is diagonal.

If S_X is shift invariant, then there is no reason for an optimal Φ to process the data in different ways for different shifts. Hence many of the basis functions of Φ will be shifts of each other. However not all shifts will be present as, within each ONB that forms Φ , there is a need for orthogonality between basis functions. For example within a discrete wavelet transform (DWT), the $n/2$ basis functions at level 1 comprise shifts that are 2 samples apart, the $n/4$ functions at level 2 are 4 samples apart, the $n/8$ functions at level 3 are 8 samples apart, and so on. The Fourier transform is well known for its ability to diagonalize shift-invariant processes, such as convolution, but it is also known for its poor localization performance, so it is usually not a good candidate for sparse signal analysis, unless it is used in its windowed and overlapped (short-time) form, when it also resembles a filter bank.

The diagonal terms in S_Y correspond to the energy of each transform-domain component, averaged over all p samples in Y . Hence for data with shift-invariant statistics and for groups of basis functions that are shifts of each other (e.g. all from within a given wavelet subband), the coefficient energies should all be the same (given sufficient training data, i.e. sufficiently large p). Thus the diagonal terms in S_Y should also be shift invariant, when they correspond to basis functions that are just shifts of each other.

Let us consider the DWT as an example of an ONB which is known to be good for a class of signals that are




Figure 2. Matrices S_X when Φ represents either a DWT (a to d) or a DTCWT (e to h), and S_Y contains unit coefficients in one wavelet level only (a,e = level 2 only; b,f = level 3 only; c,g = level 4 only; d,h = scaling functions at level 4 only).

piecewise smooth (energy only at low frequencies) but contain a relatively sparse number of step discontinuities (energy decays as $1/\omega^2$). Such signals have a power spectrum that is approximately of the form

$$P_{xx}(\omega) = \frac{2P_0T_0}{1 + (\omega T_0)^2} \quad (23)$$

where P_0 is the total power and T_0 is the time constant of the process. The autocorrelation function is the inverse Fourier transform of P_{xx} and is given by

$$r_{xx}(\tau) = P_0 \exp\left(-\frac{|\tau|}{T_0}\right) \quad (24)$$

Hence the autocorrelation matrix S_X should have a Toeplitz form in which the terms decay exponentially away from the main diagonal as $r_{xx}(\tau)$, where τ is the horizontal or vertical distance from the main diagonal.

In the wavelet domain, the energies of the wavelet coefficients of a spectrally flat (white noise) process will be constant across all subbands (since the basis functions are all normalized to unit energy). Therefore, for a process of the form given above, the energies will decay by a factor of 4 for each octave increase in frequency above the corner frequency, $1/T_0$. Hence the diagonal terms of S_Y should be proportional to [256 64 16 4 1] for the 5 subbands of a 4-level DWT, arranged in order of increasing centre frequency, and assuming that the corner frequency coincides with the upper edge of the lowest band (the level-4 scaling function). The off-diagonal components of S_Y should be zero if the transform fully decorrelates the signal, and so it should be of the form shown as an image in figure 1(a), when $n = 128$ (black pixels are more positive).

Applying the inverse DWT to both the rows and columns of S_Y , so as to obtain $S_X = \Phi S_Y \Phi^H$, produces figure 1(b). It is clear that this corresponds approximately to the decaying exponential model of a piecewise smooth process with step discontinuities, but it is also apparent that there are significant fluctuations along the diagonal direction, which should ideally not be present. These fluctuations are produced by the poor shift-invariance^{7,8} of the DWT and are the result of aliasing caused by the down-sampling processes in the DWT. This problem is illustrated in more detail in figure 2(a to d) which shows separately the components due to each subband for levels 2, 3 and 4 wavelets and for level 4 scaling functions. Hence we see that the DWT is not fully optimized for this shift-invariant process.

The dual-tree complex wavelet transform (DTCWT)^{8,9} is known to be an effective way to produce an approximately shift-invariant form of the DWT, by employing a second tree of wavelet filters in parallel with the usual DWT tree of filters. By designing the lowpass filters of the second tree to have a half-sample delay with respect to their counterparts in the first tree, the wavelet basis functions of the second tree become the approximate Hilbert transforms of those in the first tree.¹⁰ Hence, with orthonormal wavelet filters in both trees, the DTCWT forms a PTF that is a DONB, where Φ_1 represents the ONB that is due to the first tree and Φ_2 the ONB due to the second tree.

Figure 1(c) shows the S_Y matrix for a 4-level DTCWT, with $n = 128$ and $m = 256$. For clarity we have interleaved the coefficients from the two trees of the DTCWT so that pairs of equivalent coefficients from the two trees are adjacent along the diagonal of S_Y . We have used the same scaling as before, by [256 64 16 4 1], to synthesize the spectral energies of the process.

Figure 1(d) shows the S_X matrix that results from applying the DONB Φ , which represents the DTCWT, to the rows and columns of S_Y . It is clear that S_X is now very close to being shift invariant (ignoring end effects at the corners of S_X) and that it is a reasonable approximation to the exponential decay that we desire. The rate of decay may be controlled by using more (or fewer) wavelet levels and by adjusting the gain of the lowest frequency band. Figure 2(e to h) show the components due to each subband of the DTCWT separately, as was done for the DWT on the left of the figure. The improvement in shift invariance for each subband is clearly visible.

sample	Correlation values $\times 100$							
	level 1a	level 1b	level 2a	level 2b	level 3a	level 3b	level 4a	level 4b
-4	-0.00	-0.05	0.00	-0.77	0.00	1.00	0.00	2.79
-3	0.00	-0.04	-0.00	2.75	-0.00	-9.85	0.00	1.25
-2	-0.00	-0.24	0.00	-7.65	0.00	24.53	0.00	-0.92
-1	0.00	5.41	-0.00	43.38	0.00	-35.67	-0.00	0.51
0	0.00	34.76	100.00	0.01	-0.00	35.62	0.00	-1.28
1	-0.00	28.07	-0.00	-43.84	-0.00	-24.44	-0.00	-2.55
2	0.00	1.07	0.00	8.03	0.00	9.74	0.00	1.73
3	0.00	0.25	-0.00	-1.95	0.00	-0.80	0.00	2.22
4	-0.00	0.07	0.00	-0.05	0.00	-0.57	-0.00	-0.91

Table 1. Correlation between a level 2 basis function from tree a of the DTCWT, and all other nearby basis functions at levels 1 to 4 and trees a and b (ONBs Φ_1 and Φ_2). The self-correlation is shown in bold.

6. HILBERT PAIRS OF BASES

The DTCWT, considered above, is a special case of a family of DONBs, known as *Hilbert pairs of bases* (the short-time Fourier transform with 2:1 window overlap is also of this form). They have the property that each basis function in Φ_2 corresponds to the (approximate) Hilbert transform of an equivalent basis function in Φ_1 . This has several key advantages:

- The magnitude of the frequency response of a basis function is unchanged by taking its Hilbert transform (which simply shifts the phase of all frequency components by $\pm 90^\circ$). Hence the time-frequency properties of Φ_1 are preserved in Φ_2 . This helps to overcome the problem of simultaneous diagonalization of S_X by both Φ_1 and Φ_2 , encountered in equation (13).
- The $\pm 90^\circ$ phase shifts, mean that each Hilbert pair of basis functions are mutually orthogonal. This helps to produce low correlations between Φ_1 and Φ_2 in general, while still preserving similar signal decorrelation properties, and addresses the need to minimize the energy in the off-diagonal matrices of equation (13).
- The ability to separate positive from negative frequency components in the transform domain, which arises when the Hilbert pairs are combined as real and imaginary parts of complex basis functions, produces the shift invariant property that we have seen above, and also produces basis functions with good orientation selectivity⁹ when the transforms are extended into higher dimensional spaces (e.g. 2-D images and 3-D volumes).

To illustrate the decorrelation properties of Hilbert pairs of wavelet bases, we show in table 1 the correlation values between a basis function at level 2 from tree a of the DTCWT (Φ_1) and all nearby basis functions from both tree a and tree b (Φ_2) at levels 1 to 4. For this data and for that in the previous figures, 14-tap Q-shift orthonormal filters⁸ were used in the DTCWT. The correlation scores are scaled up by 100 and so are percentages of maximum correlation. Note that there is zero correlation between the input basis function (at sample 0, level 2a) and all the other functions from tree a at any level, as would be expected for an ONB. There is also approximately zero correlation (0.01%) between the input and the tree b basis function at the same level and sample position, indicating the very low correlation within a Hilbert pair. The worst correlations are of about 44% between the input and the two samples from tree b on either side of the zero-sample position. Correlations of up to 36% also exist across scale to levels 1 and 3 of tree b . There is less than 3% correlation from the level-2 input signal to level 4 basis functions.

Table 2 shows that very similar results occur when the input is a basis function from level 3, rather than level 2. This similarity exists at all levels beyond 2. At level 1, the Hilbert pair properties are not so accurate and the maximum correlations rise to 61%, instead of 44%.

sample	Correlation values $\times 100$							
	level 1a	level 1b	level 2a	level 2b	level 3a	level 3b	level 4a	level 4b
-4	0.00	0.00	0.00	0.08	0.00	-0.79	0.00	1.06
-3	0.00	-0.01	0.00	-0.07	-0.00	2.79	0.00	-10.01
-2	-0.00	-0.04	0.00	0.46	0.00	-8.27	0.00	24.75
-1	0.00	-0.06	0.00	-9.85	0.00	44.06	0.00	-35.80
0	-0.00	0.23	0.00	-35.67	100.00	0.38	-0.00	35.57
1	-0.00	2.69	-0.00	-24.44	-0.00	-44.37	-0.00	-24.24
2	0.00	0.42	0.00	-0.80	0.00	7.16	0.00	9.52
3	0.00	-0.01	0.00	-0.30	-0.00	-1.86	0.00	-0.67
4	-0.00	-0.00	0.00	-0.06	0.00	-0.07	-0.00	-0.60

Table 2. Correlation between a level 3 basis function from tree a of the DTCWT, and all other nearby basis functions at levels 1 to 4 and trees a and b (ONBs Φ_1 and Φ_2). The self-correlation is shown in bold.

It may seem that these worst-case correlation values are quite high, but values of this order are necessary when the basis functions are highly localised and thus have a very limited number of near-neighbours. The total energy of all the unwanted correlations has to sum to unity (100%), because Φ_1 and Φ_2 both have to make equal contributions to the final result when linear analysis is used. With a 1-D transform there are only about 6 near-neighbours (2 in the input level, and 2 each in the levels above and below), as can be seen in the tables. Hence the correlations with these 6 basis functions need to be of the order of $1/\sqrt{6} = 41\%$.

So far, we have been considering 1-dimensional (1-D) signals and transforms for simplicity. However many of the arguments extend naturally into 2-D and 3-D domains, particularly if separable 1-D processing (filtering) is assumed for efficiency. The Hilbert pair concept extends well into 2-D and 3-D, as long as pairs are introduced along each dimension in turn. This leads naturally to PTFs comprising $K = 4$ ONBs for 2-D images, and to PTFs with $K = 8$ ONBs for 3-D volumes.

Furthermore, in 2-D, it becomes desirable (for improved orientation selectivity) to combine the groups of 4 Hilbert-related coefficients by simple Haar pairwise sum and difference operations⁹ into pairs of complex values. Because the Haar transform is a very simple ONB, the energy-preserving properties of the union of ONBs are maintained and the resulting frame remains a PTF. Similarly, in 3-D, there are groups of 8 Hilbert-related coefficients, which may be combined by two stages of Haar transformation to produce groups of 4 complex values. The resulting frame again remains a PTF. The possibility of applying extra stages of ONBs in this way was referred to briefly at the end of section 2.

The correlations between basis functions in higher dimensions become more complicated to evaluate as there are spatial, scale and orientation directions for neighbours to exist in, so we are not able to go into this in detail here. However we have investigated this and the correlations are comparable in magnitude to the near-neighbour correlations for the 1-D case.

7. CONCLUSIONS

We have defined Parseval tight frames (PTFs) that are unions of orthonormal bases (ONBs) and have analyzed the properties of such frames. We have shown how the ideas that underly principal components analysis (PCA) may be applied to PTFs as well as to non-redundant systems. Although PCA is an inherently linear process in the non-redundant case, it is important that concepts of sparsity are introduced for the case of overcomplete frames. This leads to the need for an auxiliary variable \mathbf{z} in addition to the input vector \mathbf{x} to allow for the null-space component of each transform-domain vector \mathbf{y} . We have shown how the autocorrelation matrix for \mathbf{z} can be deduced from that for \mathbf{x} when a good sparsity-inducing analysis method is employed.

The specific case of a double ONB has been considered in some detail, and it has been shown how the requirements on the two ONBs Φ_1 and Φ_2 can be well met by designing the basis vectors of Φ_2 to be Hilbert

pairs with equivalent basis vectors of Φ_1 . Hence the dual-tree wavelet transform becomes an appropriate frame to use when the signal statistics are well-suited for wavelet analysis.

There remains considerable scope for extensions to these ideas and for more rigorous justification of some of the key concepts.

APPENDIX

Here we aim to answer the question: Why should S_Y be diagonal? (Or why does PCA work?)

Consider the case where Φ is a PTF, so that $\Phi\Phi^H = I_n$, and where $Y = \Phi^H X$, and X is the $n \times p$ matrix of p training vectors. Here we assume that the m -dimensional column vectors of Y are only in the range-space of Φ (i.e. they contain no null-space components).

$$\text{Mean energy of } X = \frac{1}{p} \text{trace}(X X^H) = \text{trace}(S_X)$$

(because it is the sum of the energies, or mean squares, of each of the n components of \mathbf{x}).

$$\text{Mean energy of } Y = \text{trace}(S_Y) = \frac{1}{p} \text{trace}(\Phi^H X X^H \Phi) = \text{trace}(\Phi^H S_X \Phi) = \text{trace}(S_X)$$

(because Φ is a PTF and preserves energy).

$$\text{Energy of the } S_X \text{ matrix (sum of squares of all its elements)} = \sum_i \sum_j |(S_X)_{ij}|^2 = \text{trace}(S_X S_X^H)$$

$$\text{Energy of the } S_Y \text{ matrix} = \text{trace}(S_Y S_Y^H) = \text{trace}(\Phi^H S_X \Phi \Phi^H S_X^H \Phi) = \text{trace}(\Phi^H S_X S_X^H \Phi) = \text{trace}(S_X S_X^H)$$

Hence both the trace and the energy of S_X are preserved in the conversion to S_Y by a Φ which is a PTF. Thus these parameters are all independent of the choice of Φ .

In order to maximize the variance of the diagonal elements of S_Y , consistent with their mean, $\text{trace}(S_Y)/m$, remaining constant, we must choose a PTF Φ which puts all the energy of S_Y into its diagonal terms. Maximizing the variance of the diagonal elements, while keeping their mean constant, has the effect of maximizing the ratio of the large terms to the smaller ones since they must all be squares and hence be non-negative. This in turn improves the sparsity of the \mathbf{y} decomposition.

Note that since we are in some sense maximizing the variance of the variances of the components of \mathbf{y} , while keeping the mean of their variances constant, this is equivalent to maximizing the kurtosis or fourth moment of the distribution of the components of \mathbf{y} .

If we now allow Y to contain vectors with null-space components Y_n as well as range-space components Y_r such that $Y = Y_r + Y_n = \Phi^H X + \Theta^H Z$, where Φ and Θ are both PTFs such that $\Phi\Theta^H = O$, then

$$S_Y = \frac{1}{p} Y Y^H = \frac{1}{p} (Y_r + Y_n) (Y_r^H + Y_n^H) = \frac{1}{p} (Y_r Y_r^H + Y_n Y_n^H) = S_{Y_r} + S_{Y_n} \quad (25)$$

since $Y_n Y_r^H = O_m$, the $m \times m$ zero-matrix, because the null-space and range-space components are mutually orthogonal.

Now we see that the energy of Y is given in terms of the range-space and null-space energies by

$$\text{trace}(S_Y) = \text{trace}(S_{Y_r}) + \text{trace}(S_{Y_n}) = \text{trace}(S_X) + \text{trace}(S_Z) \quad (26)$$

where S_Z is the covariance matrix of the null-space generating vectors in Z , and $S_{Y_n} = \Theta^H S_Z \Theta$.

The total energy of S_Y is given by

$$\begin{aligned} \text{trace}(S_Y S_Y^H) &= \text{trace}((S_{Y_r} + S_{Y_n})(S_{Y_r}^H + S_{Y_n}^H)) \\ &= \text{trace}(S_{Y_r} S_{Y_r}^H) + \text{trace}(S_{Y_n} S_{Y_n}^H) + \text{trace}(S_{Y_r} S_{Y_n}^H) + \text{trace}(S_{Y_n} S_{Y_r}^H) \end{aligned} \quad (27)$$

The final two terms above may be ignored because they are of the form

$$\text{trace}(S_{Y_r} S_{Y_n}^H) = \text{trace}(\Phi^H S_X \Phi \Theta^H S_Z \Theta) = \text{trace}(\Phi^H S_X O S_Z \Theta) = 0 \quad (28)$$

as the rows of Θ and Φ are mutually orthogonal, yielding $\Phi\Theta^H = O$, a $n \times (m - n)$ matrix of zeros.

Hence

$$\text{trace}(S_Y S_Y^H) = \text{trace}(S_{Y_r} S_{Y_r}^H) + \text{trace}(S_{Y_n} S_{Y_n}^H) = \text{trace}(S_X S_X^H) + \text{trace}(S_Z S_Z^H) \quad (29)$$

i.e. the energy of S_Y is the sum of the energy of S_{Y_r} and the energy of S_{Y_n} , which in turn equal the energies of S_X and S_Z , because Φ and Θ are PTFs.

If the trace of S_Z (energy of Z) and the energy of S_Z are assumed to be known (as well as these parameters for S_X), then by the same argument as given above for the pure range-space case, the greatest range of values of the diagonal elements of S_Y will be achieved when there is negligible energy in its off-diagonal terms. Hence we still wish to find the Φ and Θ which approximately diagonalize S_Y in order to obtain the greatest compaction of energy into the larger components of Y .

REFERENCES

- [1] Kovačević, J. and Chebira, A., “Life beyond bases: The advent of frames (Part I – Redundant: To be or not to be?),” *IEEE Signal Processing Magazine* **24**, 86–104 (July 2007).
- [2] Kovačević, J. and Chebira, A., “Life beyond bases: The advent of frames (Part II – Frames in action),” *IEEE Signal Processing Magazine* **24**, 115–125 (Sept 2007).
- [3] Akhiezer, N. I. and Glazman, I. M., [*Theory of Linear Operators in Hilbert Spaces*], vol. 1, Dover, New York (1933).
- [4] Mallat, S. G. and Zhang, Z., “Matching pursuits with time-frequency dictionaries,” *IEEE Trans. Signal Proc.* **41**, 3397–3415 (Dec. 1993).
- [5] Elad, M., Matalon, B., Shtok, J., and Zibulevsky, M., “A wide-angle view at iterated shrinkage algorithms,” in [*SPIE (Wavelets XII)*], (Aug 2007).
- [6] Donoho, D. L., Tsaig, Y., Drori, I., and Starck, J., “Sparse Solution of Underdetermined Linear Equations by Stagewise Orthogonal Matching Pursuit,” *Department of Statistics, Stanford University* <http://www-stat.stanford.edu/~donoho/Reports/2006/StOMP-20060403.pdf> (March 2006).
- [7] Simoncelli, E. P., Freeman, W. T., Adelson, E. H., and Heeger, D. J., “Shiftable multi-scale transforms,” *IEEE Trans. Inform. Theory* **38**, 587–607 (Mar. 1992).
- [8] Kingsbury, N. G., “Complex wavelets for shift invariant analysis and filtering of signals,” *Applied and Computational Harmonic Analysis* **10**, 234–253 (May 2001).
- [9] Selesnick, I., Baraniuk, R., and Kingsbury, N., “The Dual-Tree Complex Wavelet Transform,” *IEEE Signal Processing Magazine* **22**, 123–151 (Nov 2005).
- [10] Selesnick, I. W., “Hilbert transform pairs of wavelet bases,” *IEEE Signal Processing Letters* **8**, 170–173 (June 2001).