

3-D Directional Filter Banks and Surfacelets INVITED

Yue Lu and Minh N. Do

Department of Electrical and Computer Engineering
Coordinated Science Laboratory
University of Illinois at Urbana-Champaign, Urbana IL 61801, USA

ABSTRACT

In 1992, Bamberger and Smith proposed the directional filter bank (DFB) for an efficient directional decomposition of two-dimensional (2-D) signals. Due to the nonseparable nature of the system, extending the DFB to higher dimensions while still retaining its attractive features is a challenging and previously unsolved problem. This paper proposes a new family of filter banks, named 3DDFB, that can achieve the directional decomposition of 3-D signals with a simple and efficient tree-structured construction. The ideal passbands of the proposed 3DDFB are rectangular-based pyramids radiating out from the origin at different orientations and tiling the whole frequency space. The proposed 3DDFB achieves perfect reconstruction. Moreover, the angular resolution of the proposed 3DDFB can be iteratively refined by invoking more levels of decomposition through a simple expansion rule. We also introduce a 3-D directional multiresolution decomposition, named the *surfacelet transform*, by combining the proposed 3DDFB with the Laplacian pyramid. The 3DDFB has a redundancy factor of 3 and the surfacelet transform has a redundancy factor up to 24/7.

Keywords: Surfacelets, directional filter banks, filter design, directional decomposition, undecimated filter bank.

1. INTRODUCTION

Directional information is an important and unique feature of multidimensional signals. One possible scheme to obtain this information in two-dimensional (2-D) signals is through the directional filter bank (DFB), which was originally proposed by Bamberger and Smith¹ and subsequently improved by several authors.²⁻⁵ The DFB is efficiently implemented via an l -level tree-structured decomposition that leads to 2^l subbands with wedge-shaped frequency partitioning as shown in Figure 1. Meanwhile, the DFB is a non-redundant transform, and offers perfect reconstruction, i.e., the original signal can be exactly reconstructed from its decimated channels. The directional-selectivity and efficient structure of the DFB makes it an attractive candidate for many image processing applications.

One aim of this paper is to propose a novel DFB for three-dimensional signals. With the increasing capabilities of modern computers and imaging devices, high-resolution 3-D and even higher dimensional volumetric data are increasingly available in a wide gamut of scientific and technological disciplines ranging from biomedical sciences to extragalactic astronomy. We expect a new tool capable of providing efficient directional decomposition of 3-D signals can pave the way for new algorithms and applications in the processing, understanding, and manipulation of 3-D data of various sorts.

However, unlike the separable wavelets,⁶⁻⁸ whose multidimensional generalizations are simply the tensor products of their 1-D counterparts, the DFB has a much more involved non-separable construction. Extending the DFB to higher dimensions while still retaining its various attractive features is a challenging and, to the best of our knowledge, previously unsolved problem. In this paper, we propose a new family of filter banks, named 3DDFB, with the following properties:

Further author information:

Yue Lu: yuelu@uiuc.edu; Minh N. Do: minhdo@uiuc.edu

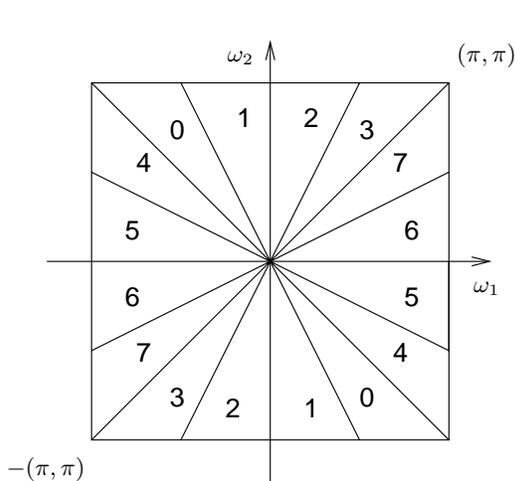


Figure 1. Frequency partitioning of the directional filter bank with 3 levels of decomposition. There are $2^3 = 8$ real wedge-shaped frequency bands.

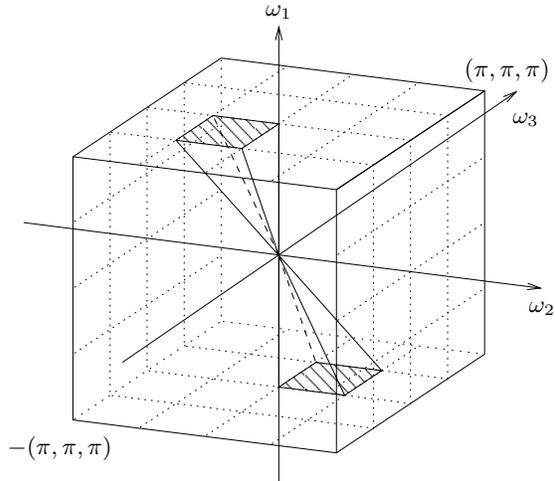


Figure 2. Frequency partitioning of the proposed 3DDFB. The ideal passbands of the component filters are rectangular-based pyramids radiating out from the origin at 3×2^l ($l \geq 0$) different orientations and tiling the whole frequency space.

1. **Directional decomposition.** The proposed 3DDFB decomposes 3-D signals into directional subbands. The component filters of the 3DDFB has ideal passbands as rectangular-based pyramids radiating out from the origin at different orientations and tiling the whole frequency space, as shown in Figure 2.
2. **Construction.** The proposed 3DDFB has a simple and efficient tree-structured implementation. The two building blocks of the 3DDFB only require simple filter design.
3. **Angular resolution.** The number of directional subbands can be doubled by iteratively invoking more levels of decomposition through a simple expansion rule. In general, there can be 3×2^l ($l \geq 0$) different directional subbands.
4. **Perfect reconstruction.** The original signal can be exactly reconstructed from its transform coefficients in the absence of noise or other processing. Actually, the proposed 3DDFB constitutes a tight frame under certain conditions.
5. **Redundancy.** The 3DDFB is a 3-times expansive system.

We are not the first to consider extending the DFB to higher dimensions. Bamberger² proposed a 3-D subband decomposition scheme implemented by applying the checkerboard filter banks separately along two orthogonal signal planes followed by a 2-D DFB decomposition on one of the planes. However, the resulting passband shapes are 3-D triangular prisms and do not correspond to a single dominant direction. Park⁴ proposed a 3-D velocity selective filter bank by applying two DFBs separately along two signal planes. The resulting frequency partitioning is similar to that of 3DDFB. However, that construction has a redundancy factor of 2^l for l -levels of decomposition, and typically $l \geq 4$. We would like to emphasize that our proposed 3DDFB has a redundant ratio of 3 in 3-D, independent of the number of decompositions.

Combining the Laplacian pyramid with the proposed 3DDFB, we propose the *surfacelet transform*, which provides a directional multiresolution decomposition of 3-D signals. A related system that can also provide directional multiresolution decomposition in multidimensions is the dual-tree complex wavelet transform (CWT).^{9–11} We would like to mention two major advantages of the proposed surfacelets over the CWT. First, in terms of redundancy, the CWT is expansive by 8 in the 3-dimensional case, while the surfacelet transform are expansive by 24/7. Second, the CWT has a fixed number of directional subbands. In contrast, the surfacelet transform can refine their angular resolution (i.e. provide more directional subbands) by invoking more levels of decomposition.

The outline of the paper is as follows. In Section 2, we give an overview of the proposed 3-D directional filter banks. In Section 3 and Section 4, we describe in detail the two building blocks, i.e., the 3-D hourglass filter bank and the 2-D iterated checkerboard filter banks, respectively. In Section 5, we introduce the 3-D *surfacelets*, as a combination of the Laplacian pyramid and the proposed 3-D directional filter banks. Section 6 concludes the paper with some discussions.

Notations: Throughout the paper, N represents the dimension of the signal. We are mainly interested in the case when $N = 3$. We use lower-case letters, e.g., $x[\mathbf{n}]$ to denote N -D discrete signals, where $\mathbf{n} \stackrel{\text{def}}{=} (n_1, n_2, \dots, n_N)^T$ is an integer vector. The z -transform of a multidimensional signal is defined as

$$X(\mathbf{z}) = \sum_{\mathbf{n} \in \mathbb{Z}^N} x[\mathbf{n}] \mathbf{z}^{-\mathbf{n}},$$

where raising an N -dimensional complex vector $\mathbf{z} \stackrel{\text{def}}{=} (z_1, \dots, z_N)^T$ to the integer vector \mathbf{n} yields $\mathbf{z}^{\mathbf{n}} = \prod_{i=1}^N z_i^{n_i}$. The discrete-time Fourier transform of a multidimensional signal is defined as

$$X(e^{j\boldsymbol{\omega}}) = \sum_{\mathbf{n} \in \mathbb{Z}^N} x[\mathbf{n}] e^{-j\boldsymbol{\omega}^T \mathbf{n}}.$$

Preliminaries: Multirate identities^{7,12} are often useful in analyzing multidimensional multirate systems. The identity for the analysis part of the filter bank is shown in Figure 3; the one for the synthesis part can be inferred similarly. Downsampling by M followed by filtering with a filter $H(\boldsymbol{\omega})$ is equivalent to filtering with the filter $H(M^T \boldsymbol{\omega})$, which is obtained by upsampling $H(\boldsymbol{\omega})$ by M , before downsampling.



Figure 3. The multidimensional multirate identity for interchanging the order of downsampling and filtering.

2. THE PROPOSED 3-D DIRECTIONAL FILTER BANK

In this section, we present an overview of the proposed 3-D directional filter bank (3DDFB). The two building blocks, i.e., the hourglass filter bank and the iterated checkerboard filter banks will be described in detail in Section 3 and Section 4, respectively.

2.1. The Hourglass Filter Bank

To obtain the first level of decomposition, we employ a three-channel undecimated filter bank shown in Figure 4. This filter bank decomposes the 3-D frequency spectrum of the input signal into three hourglass-shaped subbands, with their dominant directions aligned with the ω_1 , ω_2 , and ω_3 axes, respectively.

Despite the redundancy it brings in, the undecimated hourglass filter bank in this step offers several important advantages over a decimated filter bank:

1. As will be seen shortly, the hourglass filter bank allows subsequent levels of the 3DDFB to be implemented by two 2-D filter banks working separately along two orthogonal signal planes. This simplifies the design and implementation of the overall system.
2. Designing nonseparable 3-D filter banks with perfect reconstruction and good frequency selectivity is still a very challenging problem. In general, it is much easier to design an undecimated filter bank than a decimated one, since the former imposes a smaller set of constraints.

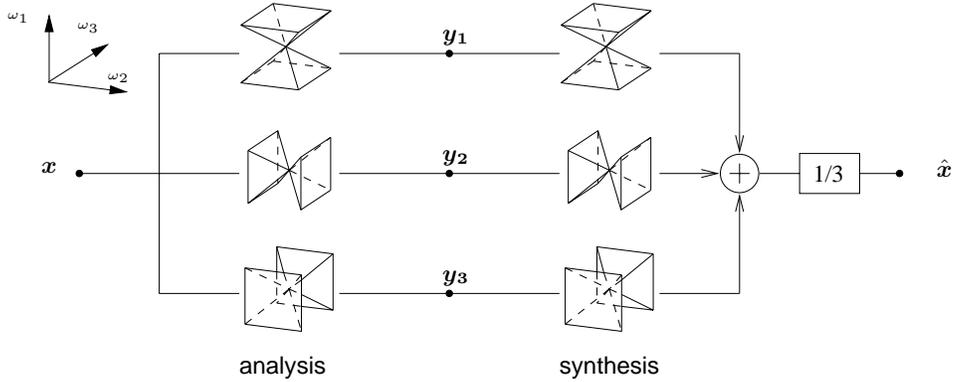


Figure 4. The first level of decomposition: a three-channel undecimated filter bank in 3-D. The ideal frequency-domain supports of the component filters are hourglass-shaped regions, with their corresponding dominant directions aligned with the ω_1 , ω_2 , and ω_3 axes, respectively.

We should make several simplifications before describing further levels of decomposition. First, we only focus on the analysis part of the proposed 3DDFB, since the synthesis part is exactly symmetric. Second, we only consider subsequent decomposition steps after the top branch of the hourglass filter bank, i.e., the y_1 subband in Figure 4 whose dominant direction is along the ω_1 axis. The decomposition after the two lower branches can be constructed by permuting the three dimensions, e.g., $(n_1, n_2, n_3) \rightarrow (n_2, n_3, n_1)$ and $(\omega_1, \omega_2, \omega_3) \rightarrow (\omega_2, \omega_3, \omega_1)$, from the corresponding channels in the top branch. This applies to both the sampling matrices and the filters used in the decomposition. Third, since we are mainly interested in the passband and stopband regions of the filters, we assume all the filters used in this section and Section 4 are ideal, i.e., the frequency response of each filter takes the values one in its passband and zero in its stopband.

2.2. The Generalized Separability of the Pyramid-Shaped Filters

We first introduce some notations for several frequently-used support regions. Figure 5(a) shows the wedge-shaped decomposition of the 3-D frequency spectrum. Since frequency supports remain the same along the ω_3 axis, this decomposition can be achieved by applying a 2-D filter bank, e.g., the DFB, along the (n_1, n_2) -plane. We use

$$W_i^{(l_1)}(\omega_1, \omega_2), \quad 0 \leq i \leq 2^{l_1} - 1,$$

to denote the ideal filter whose frequency support is on the i th wedge. The superscript (l_1) indicates that there are 2^{l_1} wedge subbands (in this case, $l_1 = 2$) oriented at angles from -45° to 45° . The frequency variables ω_1 and ω_2 specify that the 2-D filters operate along the (n_1, n_2) -plane. Similarly, we show in Figure 5(b) the wedge-shaped frequency decomposition along the (n_1, n_3) -plane. With the same notation above, we can use $W_j^{(l_2)}(\omega_1, \omega_3)$ ($0 \leq j \leq 2^{l_2} - 1$) to represent the ideal subband filters. Figure 5(c) shows the pyramid-shaped frequency decomposition. Ideally, the hourglass-shaped region is divided into $2^{l_1} \times 2^{l_2}$ ($l_1, l_2 \geq 0$) different square-based pyramids radiating out from the origin. As illustrated in Figure 5(c), each subband can be indexed by a pair (i, j) specifying the square base of the pyramid. We use

$$P_{i,j}^{(l_1, l_2)}(\omega_1, \omega_2, \omega_3), \quad 0 \leq i \leq 2^{l_1} - 1 \text{ and } 0 \leq j \leq 2^{l_2} - 1,$$

to represent the (i, j) th ideal subband filter. As a special case, the ideal hourglass-shaped filter aligned along the ω_1 axis can be denoted as $P_{0,0}^{(0,0)}(\omega_1, \omega_2, \omega_3)$.

We now start showing how to achieve this frequency decomposition in an efficient tree-structured approach.

Separability is an important concept in multidimensional signal processing. A multidimensional filter $F(\boldsymbol{\omega})$ is *separable*, when it can be written as the product of several 1-D filters, i.e.

$$F(\boldsymbol{\omega}) = \prod_{i=1}^N F_i(\omega_i).$$

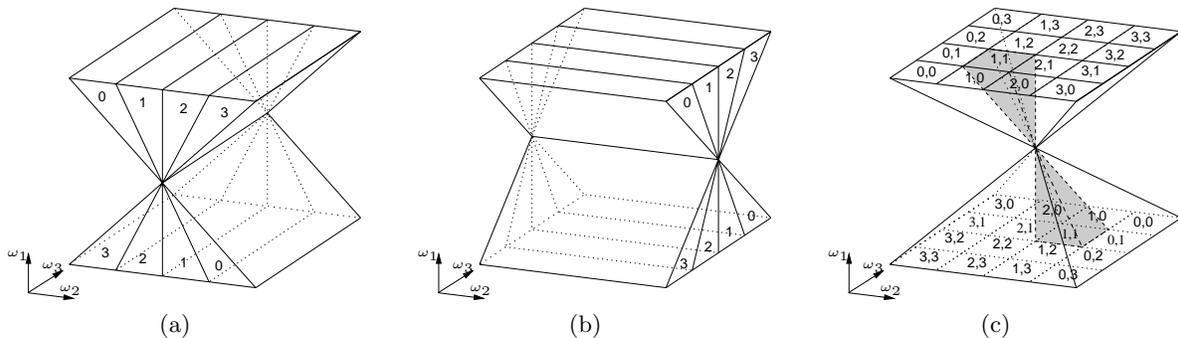


Figure 5. (a) The ideal wedge-shaped frequency support of a 2-D filter operating along the (n_1, n_2) -plane. (b) The wedge-shaped support of a 2-D filter along the (n_1, n_2) -plane. (c) The ideal pyramid-shaped frequency decomposition.

Here, we introduce the notion of *generalized separability with order K* , to describe those N -dimensional filters that can be represented as the product of several K -dimensional filters, with $0 < K < N$. A very important observation is: the pyramid filters $P_{i,j}^{(l_1, l_2)}(\omega_1, \omega_2, \omega_3)$ defined above are generalized separable with order 2.

LEMMA 2.1 (GENERALIZED SEPARABILITY).

$$P_{i,j}^{(l_1, l_2)}(\omega_1, \omega_2, \omega_3) = W_i^{(l_1)}(\omega_1, \omega_2) \cdot W_j^{(l_2)}(\omega_1, \omega_3). \quad (1)$$

for all $l_1, l_2 \geq 0$ and $0 \leq i \leq 2^{l_1} - 1$, $0 \leq j \leq 2^{l_2} - 1$.

This can be verified by simple geometric arguments. At this point, a natural question is: since the pyramid filters are the products of two wedge filters, can we achieve the pyramid-shaped frequency decomposition by applying two 2-D DFBs separately along the (n_1, n_2) and (n_1, n_3) signal planes? This idea was explored by Park,⁴ but at a high price.

Here is why. Let $W_i^{(l_1)}(\omega_1, \omega_2)$ and $W_j^{(l_2)}(\omega_1, \omega_3)$ represent two wedge-shaped filters from the two DFBs, respectively. The key problem is that the DFB is critically-sampled. The equivalent downsampling matrix for the DFB along the (n_1, n_2) -plane is⁵

$$\mathbf{M} = \begin{pmatrix} 2 & 0 & 0 \\ 0 & 2^{l_1} & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

Sampling by \mathbf{M} , in particular with a downsampling by 2 along the n_1 dimension, scrambles the wedge-shaped frequency decomposition provided by $W_i^{(l_1)}(\omega_1, \omega_3)$ as in Figure 5(a). Thus it can be easily checked that the subsequent application of $W_j^{(l_2)}(\omega_1, \omega_3)$ will not provide the desired pyramid-shaped frequency decomposition as shown in Figure 5(c).

To get rid of this problem, Park⁴ proposed to upsample and interpolate the decimated outputs of the first DFB (by a synthesis filter bank) to the original size before feeding them to the second DFB. With this step, the DFB essentially becomes an undecimated filter bank, and hence the generalized separability can be applied. However, this scheme leads to a highly redundant system. In general, that construction is 2^l -times redundant for l -levels of decomposition, and typically $l \geq 4$.

2.3. Subsequent Levels of Decomposition

For finer frequency partition, we propose a new filter bank structure, that can make use of the generalized separability property, but without the redundancy. As shown in Figure 6, we sequentially apply two 2-D filter banks after the hourglass filter, with the first one, denoted as $\mathcal{S}^{(l_1)}(n_1, n_2)$, operating along the (n_1, n_2) -plane slice by slice and the second one, $\mathcal{S}^{(l_2)}(n_1, n_3)$, along the (n_1, n_3) -plane slice by slice. $\mathcal{S}^{(l_1)}(n_1, n_2)$ has a critically-sampled binary-tree structure with l_1 ($l_1 \geq 0$) levels of decomposition, and therefore has 2^{l_1} different output

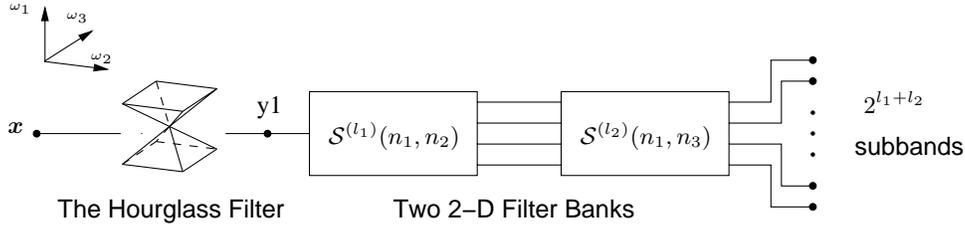


Figure 6. The proposed filter bank structure of the 3DDFB. The hourglass filter in the first level is followed by two 2-D filter banks $\mathcal{S}^{(l_1)}(n_1, n_2)$ and $\mathcal{S}^{(l_2)}(n_1, n_3)$ operating along two orthogonal signal planes.

branches. We can rewrite the tree-structured filter bank $\mathcal{S}^{(l_1)}(n_1, n_2)$ in its equivalent parallel form, where each branch is consisted of one equivalent filter followed by an equivalent downsampling matrix.

We use $S_i^{(l_1)}(\omega_1, \omega_2)$ to denote the equivalent filter of the i th branch. As suggested by its notation, the second filter bank $\mathcal{S}^{(l_2)}(n_1, n_3)$ has the same construction as $\mathcal{S}^{(l_1)}(n_1, n_2)$, but operating along a different signal plane, i.e., $(n_1, n_2) \rightarrow (n_1, n_3)$, and with a different decomposition depth, i.e. $l_1 \rightarrow l_2$. Note that $\mathcal{S}^{(l_2)}(n_1, n_3)$ is attached to every subband of $\mathcal{S}^{(l_1)}(n_1, n_2)$, so we have a total of $2^{l_1+l_2}$ output subbands. We use $S_{i,j}^{(l_1, l_2)}(\omega_1, \omega_2, \omega_3)$ to denote the equivalent filter that comes from the (i, j) th output subband.

Our goal is to have the pyramid-shaped frequency decomposition given in Figure 5(c). To achieve this, the 2-D filter bank $\mathcal{S}^{(l_1)}(n_1, n_2)$ needs to satisfy the following two requirements.

1. Equivalent Downsampling Matrices. Let \mathbf{P} denote the equivalent downsampling matrix for each subband of $\mathcal{S}^{(l_1)}(n_1, n_2)$, then we want

$$\mathbf{P} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 2^{l_1} & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (2)$$

This condition serves two purposes. First, $\mathcal{S}^{(l_1)}(n_1, n_2)$ has 2^{l_1} subbands while the determinant of \mathbf{P} is also 2^{l_1} . This ensures we have a maximally-decimated (nonredundant) filter bank. The other (more important) feature of \mathbf{P} is that it is a diagonal matrix with the first and third diagonal elements being one. In calculating the total equivalent filter, the second filter bank $\mathcal{S}_j^{(l_2)}(\omega_1, \omega_3)$ is only upsampled by an identity matrix and hence will not be changed. We have

$$S_{i,j}^{(l_1, l_2)}(\omega_1, \omega_2, \omega_3) = P_{0,0}^{(0,0)}(\omega_1, \omega_2, \omega_3) \cdot S_i^{(l_1)}(\omega_1, \omega_2) \cdot S_j^{(l_2)}(\omega_1, \omega_3), \quad (3)$$

for all $l_1, l_2 \geq 0$ and $0 \leq i \leq 2^{l_1} - 1$, $0 \leq j \leq 2^{l_2} - 1$.

2. Equivalent Filters. Using Lemma 2.1, we can decompose the hourglass filter in (3) as the product of two “virtual” wedge filters:

$$P_{0,0}^{(0,0)}(\omega_1, \omega_2, \omega_3) = W_0^{(0)}(\omega_1, \omega_2) \cdot W_0^{(0)}(\omega_1, \omega_3).$$

Now (3) can be rewritten as

$$S_{i,j}^{(l_1, l_2)}(\omega_1, \omega_2, \omega_3) = \left(W_0^{(0)}(\omega_1, \omega_2) \cdot S_i^{(l_1)}(\omega_1, \omega_2) \right) \cdot \left(W_0^{(0)}(\omega_1, \omega_3) \cdot S_j^{(l_2)}(\omega_1, \omega_3) \right). \quad (4)$$

Recall that our goal is to have $S_{i,j}^{(l_1, l_2)}(\omega_1, \omega_2, \omega_3)$ being the pyramid filters $P_{i,j}^{(l_1, l_2)}(\omega_1, \omega_2, \omega_3)$. By comparing (4) with (1), we reach the second important requirement:

$$W_i^{(l_1)}(\omega_1, \omega_2) = W_0^{(0)}(\omega_1, \omega_2) \cdot S_i^{(l_1)}(\omega_1, \omega_2). \quad (5)$$

As shown in Section 4, the filter bank $\mathcal{S}^{(l)}(n_1, n_2)$ that satisfy both (2) and (5) turns out to be an iterated interconnection of a checkerboard filter bank and some 2-D resampling matrices

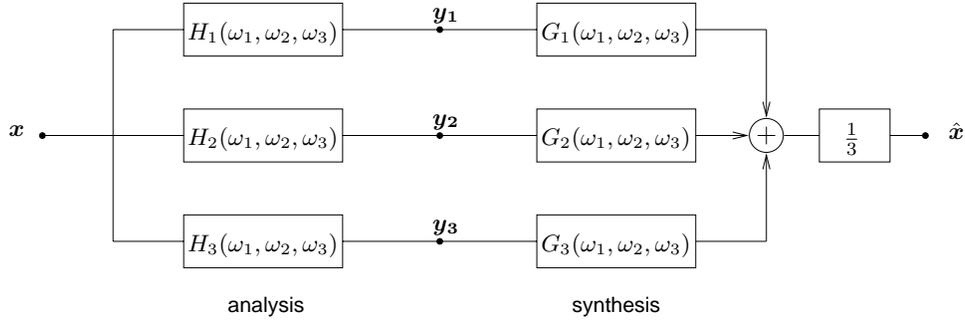


Figure 7. The undecimated hourglass filter bank in 3-D.

3. DESIGNING THE HOURGLASS FILTER BANKS IN 3-D

Our goal here is to design the 3-D undecimated filter bank in Figure 4 with perfect reconstruction and the desired hourglass-shaped frequency decomposition. Generally speaking, designing three and higher-dimensional filter banks is a very challenging task with few ready-to-use tools available.

In this paper, we propose a novel design based on frequency-domain techniques. As shown in Figure 7, we use $H_i(\omega_1, \omega_2, \omega_3)$ and $G_i(\omega_1, \omega_2, \omega_3)$ for $i = 1, 2, 3$ to represent the three analysis and synthesis filters in the hourglass filter bank, respectively. As the first step of simplification, we assume the three analysis filters are rotational-symmetric to each other, i.e.,

$$H_2(\omega_1, \omega_2, \omega_3) = H_1(\omega_3, \omega_1, \omega_2) \quad \text{and} \quad H_3(\omega_1, \omega_2, \omega_3) = H_1(\omega_2, \omega_3, \omega_1).$$

The same constraint also applies to the synthesis filters. Meanwhile, if the filter bank implements a tight frame expansion, we need the synthesis filters to be the time-reversed versions of the corresponding analysis filters, i.e.,

$$G_i(\boldsymbol{\omega}) = H_i(-\boldsymbol{\omega}) = H_i(\boldsymbol{\omega}),$$

for $i = 1, 2, 3$, where the second equality comes from the symmetry in the ideal frequency responses of $H_i(\boldsymbol{\omega})$. Combining the above two constraints, we get the condition for perfect reconstruction as

$$H_1^2(\omega_1, \omega_2, \omega_3) + H_1^2(\omega_3, \omega_1, \omega_2) + H_1^2(\omega_2, \omega_3, \omega_1) = 3. \quad (6)$$

Inspired by the work by Feilner et al.¹³ on 2-D quincunx wavelets, we propose a novel construction, in which we let

$$H_1(\omega_1, \omega_2, \omega_3) = \sqrt{\frac{3 \cdot F(\omega_1, \omega_2, \omega_3)^\lambda}{F(\omega_1, \omega_2, \omega_3)^\lambda + F(\omega_3, \omega_1, \omega_2)^\lambda + F(\omega_2, \omega_3, \omega_1)^\lambda}},$$

where $\lambda > 0$ and $F(\omega_1, \omega_2, \omega_3) > 0$ is a positive and 2π periodic function of ω_1, ω_2 and ω_3 .

We can verify that the perfect reconstruction condition in (6) is satisfied by arbitrary choices of λ and $F(\omega_1, \omega_2, \omega_3)$. To control the frequency responses of the analysis filters so that they approximate the desired hourglass shape, we let

$$F(\omega_1, \omega_2, \omega_3) = (1 + P(\cos(\omega_1), \cos(\omega_2))) (1 + P(\cos(\omega_1), \cos(\omega_3))),$$

where $P(\cdot, \cdot)$ is a bivariate polynomial such that $P(\cos(\omega_1), \cos(\omega_2))$ approximately takes the value +1 in the dark region in Figure 8 and the value -1 in the white region.

There can be many ways in designing the polynomial $P(\cdot, \cdot)$. In our experiment, we employ the windowing method proposed by Tay and Kingsbury¹⁴ and choose $\lambda = 4$. Figure 9 shows the isosurface of the frequency response of one analysis filter. We can see that the frequency response approximates the ideal hourglass shape fairly well. Note that the responses of other filters are rotational-symmetric to this one.

The main objection that can be made to this construction is that the resulting filters do not have finite impulse responses (FIR). Despite this unfavorable property, the proposed design offers many desirable features, including:

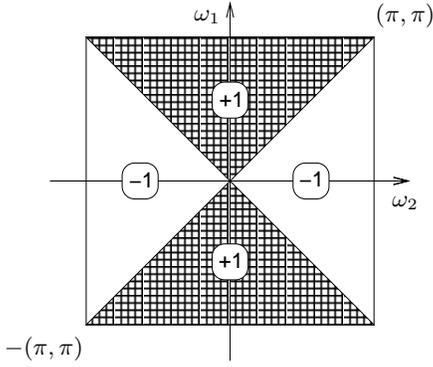


Figure 8. $P(\cos(\omega_1), \cos(\omega_2))$ approximately takes the value $+1$ in the dark region and the value -1 in the white region.

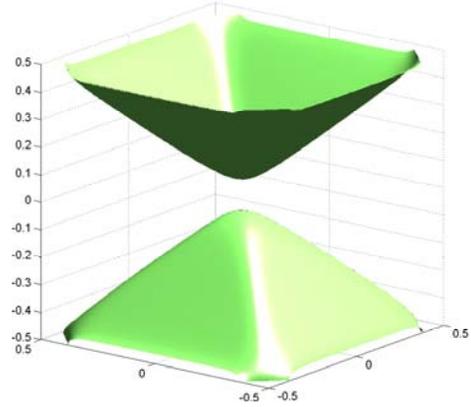


Figure 9. The frequency response of one analysis filter designed by the proposed frequency-domain method. The responses of other filters are rotational-symmetric to this one.

1. Perfect reconstruction;
2. The analysis and synthesis filters approximate the ideal hourglass-shaped frequency responses;
3. The three analysis filters are rotational-symmetric to each other;
4. The filter bank constitutes a tight frame of redundancy 3;
5. Simple and efficient frequency-domain implementation.

4. CONSTRUCTING THE ITERATED CHECKERBOARD FILTER BANKS

In this section we construct the filter bank $\mathcal{S}^{(l)}(n_1, n_2)$ that satisfies requirements (2) and (5).

When $l = 0$, the filter bank $\mathcal{S}^{(0)}(n_1, n_2)$ is simply the identity transform. We need to consider this degenerate case, since sometimes we just want to decompose the 3-D hourglass support along only one, e.g. the ω_3 , direction.

When $l = 1$, the filter bank $\mathcal{S}^{(1)}(n_1, n_2)$ is a two-channel 2-D filter bank with a checkerboard-shaped frequency partition, as illustrated in Figure 10. Note that we need to attach two resampling operations, denoted as \mathbf{R}_0 and \mathbf{R}_1 , to channel 0 and channel 1, respectively.

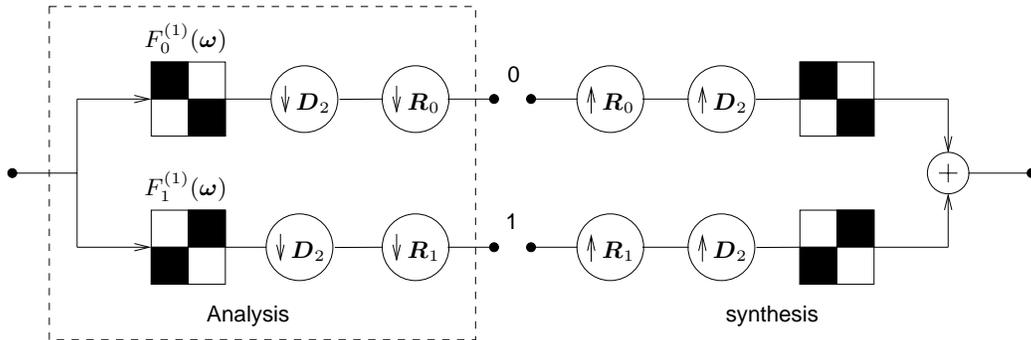


Figure 10. The two-channel 2-D checkerboard filter bank with resampling. The dark regions represent the ideal passband.

The sampling matrices in Figure 10 are defined as

$$D_2 = \begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix}, \quad \mathbf{R}_0 = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad \mathbf{R}_1 = \begin{pmatrix} 1 & -1 \\ 0 & 1 \end{pmatrix}.$$

When $l > 1$, the filter bank $\mathcal{S}^{(l)}(n_1, n_2)$ is an iterative expansion of the resampled checkerboard filter banks. Specifically, we build an l -level binary tree (in the analysis part) by recursively attaching a copy of the diagram contents enclosed by the dashed rectangle in Figure 10 to every output channels from the previous level.

We can index the channels of the $\mathcal{S}^{(l)}(n_1, n_2)$ from top to bottom with the integers from 0 to $2^l - 1$. Associated with each channel indexed by k ($0 \leq k \leq 2^l - 1$) is a sequence of path types (where a type is either 0 for the upper branch or 1 for the lower branch, as shown in Figure 10) (t_1, t_2, \dots, t_l) from the first level leading to that channel. According to the expanding rule, $(t_1 - 1, t_2 - 1, \dots, t_l - 1)$ is the binary representation of k , or

$$k = \sum_{i=1}^l t_i 2^{l-i}.$$

With this path type, the sequence of filtering and downsampling for channel k can be written as

$$\rightarrow \boxed{F_{t_1}^{(1)}} \rightarrow (\downarrow \mathbf{D}_2 \cdot \mathbf{R}_{t_1}) \rightarrow \boxed{F_{t_2}^{(1)}} \rightarrow (\downarrow \mathbf{D}_2 \cdot \mathbf{R}_{t_2}) \rightarrow \dots \rightarrow \boxed{F_{t_l}^{(1)}} \rightarrow (\downarrow \mathbf{D}_2 \cdot \mathbf{R}_{t_l}).$$

From this, using the multirate identities recursively, we can transform the analysis side of the channel k ($0 \leq k < 2^l$) of the $\mathcal{S}^{(l)}(n_1, n_2)$ into a single filtering with the equivalent filter $F_k^{(l)}(\boldsymbol{\omega})$ followed by downsampling by the overall sampling matrix $\mathbf{M}_k^{(l)}$, where

$$\mathbf{M}_k^{(l)} = \prod_{i=1}^l (\mathbf{D}_2 \cdot \mathbf{R}_{t_i}), \quad (7)$$

$$F_k^{(l)}(\boldsymbol{\omega}) = F_{t_1}^{(1)}(\boldsymbol{\omega}) \prod_{n=2}^l F_{t_n}^{(1)} \left((\mathbf{M}_{\lfloor k/2^{l+1-n} \rfloor}^{(n-1)})^T \boldsymbol{\omega} \right). \quad (8)$$

The matrix $\mathbf{M}_{\lfloor k/2^{l+1-n} \rfloor}^{(n-1)}$ ($n = 2, \dots, l$) in (8) is understood as the partial product of the overall sampling matrix in (7), i.e.

$$\mathbf{M}_{\lfloor k/2^{l+1-n} \rfloor}^{(n-1)} = \prod_{i=1}^{n-1} (\mathbf{D}_2 \cdot \mathbf{R}_{t_i}).$$

PROPOSITION 1. *The overall sampling matrix for the k -th ($0 \leq k < 2^l$) channel in the l -level filter bank $\mathcal{S}^{(l)}(n_1, n_2)$ is*

$$\mathbf{M}_k^{(l)} = \mathbf{D}_2^l \cdot \mathbf{R}_0^{2^l - 1 - 2k}. \quad (9)$$

To satisfy the sampling matrix condition given in (2), we need to attach a resampling matrix to each of the 2^l output channels of the $\mathcal{S}^{(l)}(n_1, n_2)$. The resampling matrix for the k th channel is defined as

$$\widetilde{\mathbf{M}}_k^{(l)} = \mathbf{R}_1^{2^l - 1 - 2k},$$

so that the overall equivalent downsampling matrix is

$$\mathbf{P} = \mathbf{M}_k^{(l)} \cdot \widetilde{\mathbf{M}}_k^{(l)} = \mathbf{D}_2^l,$$

which satisfies the condition in (2).

Now we need to verify that the proposed $\mathcal{S}^{(l)}(n_1, n_2)$ also satisfies the equivalent filter condition given in (5). The special case when $l = 1$ is illustrated in Figure 11, where a wedge-shaped support $W_0^{(0)}(\omega_1, \omega_2)$ (Figure 11(a)) is divided by the checkerboard filter (Figure 11(b)) provided by $\mathcal{S}^{(1)}(n_1, n_2)$. The result is a “thinner” wedge

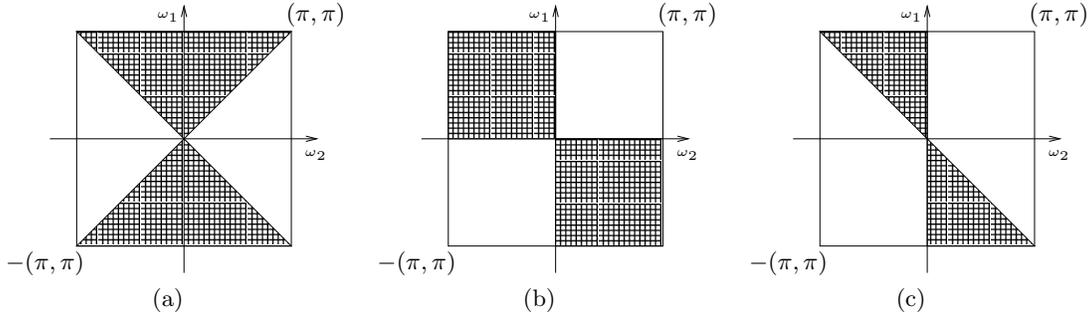


Figure 11. Verifying the equivalent filter condition in (5). (a) The ideal frequency support of $W_0^{(0)}(\omega_1, \omega_2)$. (b) The ideal checkerboard-shaped frequency support of $F_0^{(1)}(\omega_1, \omega_2)$. (c) Multiplying the supports in (a) and (b), we get the desired wedge-shaped frequency support of $W_0^{(1)}(\omega_1, \omega_2)$.

support $W_0^{(1)}(\omega_1, \omega_2)$ shown in Figure 11(c). In general, for $l > 1$, we can show that the condition (5) still holds, as stated in the following proposition.

PROPOSITION 2. Assume the l -level filter bank $\mathcal{S}^{(l)}(n_1, n_2)$ uses ideal filters with binary-valued frequency responses. The equivalent filter $F_k^{(l)}(\omega)$ for the k -th channel, $0 \leq k < 2^l$, satisfies the equivalent filter condition, i.e.,

$$W_k^{(l)}(\omega) = W_0^{(0)}(\omega) \cdot F_k^{(l)}(\omega),$$

for all $l \geq 0$.

5. THE SURFACELET TRANSFORM

Similar to the idea in the contourlet transform, we combine the proposed 3-D directional filter bank with the Laplacian pyramid, and construct the 3-D *surfacelet transform*. Figure 12 shows the block diagram of the transform. The surfacelets offer a directional multiresolution decomposition of 3-D signals with a redundancy ratio of up to 24/7. We show in Figure 13(a) the frequency support of a 3DDFB subband using non-ideal filters. Figure 13(b) and Figure 13(c) are some surfacelets in the frequency and spatial domain, respectively.

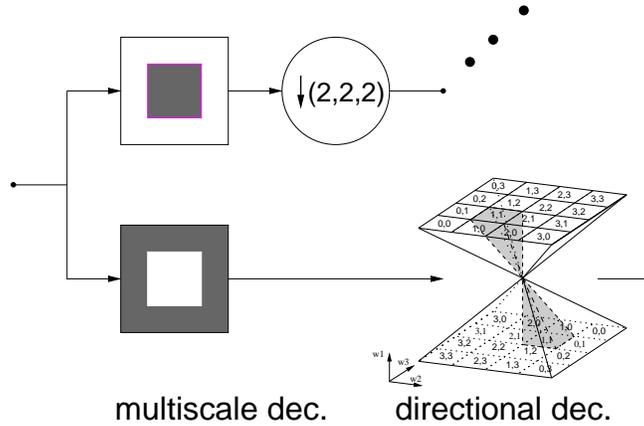


Figure 12. The block diagram of the proposed surfacelet transform. The 3-D Laplacian pyramid and the 3-D DFB are combined to form a directional multiresolution decomposition of 3-D signals.

6. CONCLUSION AND DISCUSSIONS

In this paper, we proposed a novel 3-D directional filter bank. We showed that the directional decomposition of 3-D signals is possible by employing an hourglass-shaped undecimated filter bank together with the iterated

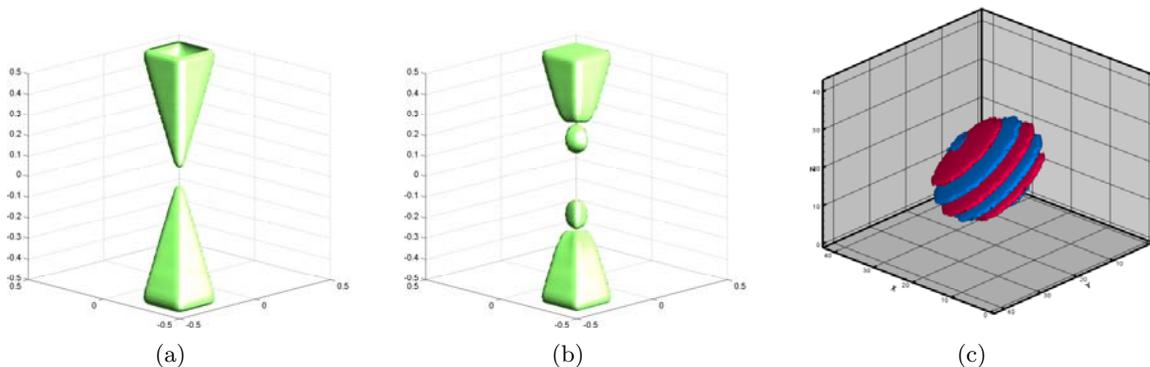


Figure 13. (a) The frequency support of a 3DDFB subband using non-ideal filters. (b) The isosurface of one surfacelet in the frequency domain. (c) The isosurface of one surfacelet in the spatial domain.

checkerboard filter banks. The proposed filter bank has perfect reconstruction, and an efficient tree-structured implementation. Combining the Laplacian pyramid with the proposed 3DDFB, we constructed the surfacelets, for a 3-D directional multiresolution analysis.

All the constructions described in this paper can be easily generalized to arbitrary N dimensional cases for $N \geq 2$. The result is the N -dimensional surfacelets that can be used to capture singularities living on $(N - 1)$ -dimensional hyper-surfaces. More details on this generalization and numerical experiments will be given in a forthcoming paper.

REFERENCES

1. R. H. Bamberg and M. J. T. Smith, "A filter bank for the directional decomposition of images: theory and design," *IEEE Trans. Signal Proc.* **40**, pp. 882–893, April 1992.
2. R. H. Bamberg, "New results on two and three dimensional directional filter banks," in *Twenty-Seventh Asilomar Conf. on Signals, Systems and Computers*, **2**, pp. 1286–1290, 1993.
3. S. Park, M. J. T. Smith, and R. M. Mersereau, "A new directional filterbank for image analysis and classification," in *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Proc.*, pp. 1417–1420, 1999.
4. S. Park, *New Directional Filter Banks and Their Applications in Image Processing*. PhD thesis, Georgia Institute of Technology, 1999.
5. M. N. Do, *Directional Multiresolution Image Representations*. PhD thesis, Swiss Federal Institute of Technology, Lausanne, Switzerland, December 2001, <http://www.ifp.uiuc.edu/~minhdo/publications/thesis.pdf>.
6. I. Daubechies, *Ten Lectures on Wavelets*, SIAM, Philadelphia, PA, 1992. Notes from the 1990 CBMS-NSF Conference on Wavelets and Applications at Lowell, MA.
7. M. Vetterli and J. Kovačević, *Wavelets and Subband Coding*, Prentice-Hall, 1995.
8. S. Mallat, *A Wavelet Tour of Signal Processing*, Academic Press, 1998.
9. N. Kingsbury, "Complex wavelets for shift invariant analysis and filtering of signals," *Journal of Appl. and Comput. Harmonic Analysis* **10**, pp. 234–253, 2001.
10. I. W. Selesnick, "The double-density dual-tree DWT," *IEEE Trans. Signal Proc.* **52**, pp. 1304–1314, May 2004.
11. R. G. Baraniuk, N. Kingsbury, and I. W. Selesnick, "The dual-tree complex wavelet transform - a coherent framework for multiscale signal and image processing," *IEEE SP Mag.*, submitted, 2004.
12. P. P. Vaidyanathan, *Multirate Systems and Filter Banks*, Prentice Hall, 1993.
13. M. Feilner, D. V. D. Ville, and M. Unser, "An orthogonal family of quincunx wavelets with continuously adjustable order," *IEEE Trans. Image Proc.* **14**(4), p. 499, 2005.
14. D. B. H. Tay and N. Kingsbury, "Flexible design of multidimensional perfect reconstruction FIR 2-band filters using transformations of variables," *IEEE Trans. Image Proc.* **2**, pp. 466–480, 1993.