

TFDWT: Fast Discrete Wavelet Transform TensorFlow Layers

Kishore K. Tarafdar¹ and Vikram M. Gadre¹

¹ Department of Electrical Engineering, Indian Institute of Technology Bombay, India

DOI: [10.xxxxxx/draft](https://doi.org/10.xxxxxx/draft)

Software

- [Review](#)
- [Repository](#)
- [Archive](#)

Editor: [✉](#)

Submitted: 31 March 2025

Published: unpublished

License

Authors of papers retain copyright and release the work under a Creative Commons Attribution 4.0 International License ([CC BY 4.0](https://creativecommons.org/licenses/by/4.0/)).

1 Summary

TFDWT is an open-source Python library that allows the construction of TensorFlow Layers for Fast Discrete Wavelet Transform (DWT) and Inverse Discrete Wavelet Transform (IDWT) in end-to-end backpropagation learning networks. These layers facilitate the construction of multilevel DWT filter banks and Wavelet Packet Transform (WPT) filter banks for a spatial-frequency representation of the inputs and features in shallow or deep networks. A multiresolution signal representation using a multi-rate discrete wavelet system creates enriched joint natural-frequency representations. The discrete wavelet system partitions the frequency plane into subbands using orthogonal dilated and translated lowpass scaling and highpass wavelet function. A realization of a fast discrete wavelet system is a two-band perfect reconstruction multi-rate filter bank with FIR filters corresponding to the impulse responses of the scaling and wavelet function with downsampling and upsampling operations. A filter bank for a higher dimensional input is a seamless extension by successive separable circular convolutions across each independent axis. The command `pip install TFDWT` installs the latest version of the package.

2 Statement of need

In machine or deep learning, an efficient multiresolution representation of data often helps to build economical and explainable models. The Wavelet toolbox (Misiti et al., 1996) by MathWorks is a proprietary software that has served the requirements for D-dimensional wavelet transforms in the MATLAB environment for a few decades. Several open-source packages are now available for 1D and 2D DWT in Python. Pywavelets (Lee et al., 2019) is a D-dimensional wavelet transform library in Python that works with Numpy (Harris et al., 2020) arrays. However, it is challenging to directly use Pywavelets with the symbolic tensors in TensorFlow (Abadi et al., 2015) layers and CUDA (Fatica, 2008). WaveTF (Versaci, 2021) is a solution for constructing 1D and 2D DWT layers in TensorFlow but is limited to only Haar and Db2 wavelets. The package tensorflow-wavelets (Leiderman, 2025) supports multiple wavelets, but it has a minor bug in perfect reconstruction due to the padding and boundary effects in processing the finite-length inputs. In Pytorch (Imambi et al., 2021), the pytorch-wavelets (Cotter, 2023) package allows the construction of 1D and 2D DWT layers. However, there are limited libraries for 3D and higher dimensional transforms with a wide range of wavelet families for Graphics Processing Unit (GPU) computations.

For a D-dimensional wavelet $\psi \in L^2(\mathbb{R})^D$, a discrete wavelet system defined by $\{\psi_{m,p} : m \in \mathbb{Z}, p \in \mathbb{Z}^D, D \in \mathbb{N}\}$ forms an orthonormal basis, where $\psi_{m,p}(x) := 2^m \psi(2^m x - p)$. Then, by definition the DWT of $x \in \mathbb{Z}^D$ is $x \mapsto (\langle x, \psi_{m,p} \rangle)_{m,p}$, where m is the dilation parameter and p is the shift or translation parameter. The TFDWT Python package is a simple standalone DWT and IDWT library with minimal dependencies that allow computation with symbolic tensors and CUDA. This release supports up to 3D forward and inverse transforms

with various orthogonal and biorthogonal wavelet families. A seamless package upgrade for higher dimensional DWT is possible by separable transforms of the independent axes. The boundary effects are taken care of with cyclic convolutions instead of padding. The package supports orthogonal and biorthogonal wavelets of different families having impulse responses of diverse lengths. In this paper, we defined the platform-independent, underlying mathematics for realizing fast D-dimensional DWT and IDWT layers with filter bank structures having FIR filters, downsamplers and upsamplers. Although our realization of the discrete wavelet system is in TensorFlow 2, a seamless reproduction of the computations is possible in other deep learning frameworks.

3 Discrete wavelet system for sequences

A discrete wavelet system contains a pair of quadrature mirror filters with a lowpass scaling function and a highpass mother wavelet. A discrete wavelet system with a continuous one-dimensional ($D = 1$) mother wavelet $\psi \in L^2(\mathbb{R})$ is realized by using the impulse responses of the scaling function and the wavelet as Finite Impulse Response (FIR) filters $g[n]$ and $h[n]$. Figure 1 shows wavelets and scaling functions of the analysis and synthesis bank of bior3.1 wavelet and their corresponding impulse responses are shown in Figure 2. These FIR filters are the building blocks of a two-band perfect reconstruction filter bank for realizing fast discrete wavelet systems. Figure 3 shows a two-band perfect reconstruction filter bank that operates on one-dimensional inputs, i.e., sequences in $l^2(\mathbb{Z})$. The analysis and synthesis bank in orthogonal wavelet systems have identical lowpass and highpass FIR filters but differ in biorthogonal wavelet systems. In biorthogonal wavelet filter banks, $\tilde{g}[n]$ and $\tilde{h}[n]$ are the lowpass and highpass filters of the synthesis bank. The only difference in the biorthogonal families like bior and rbio is the interchange of the analysis and synthesis scaling and wavelets functions, for example, in bior3.1 and rbio3.1.

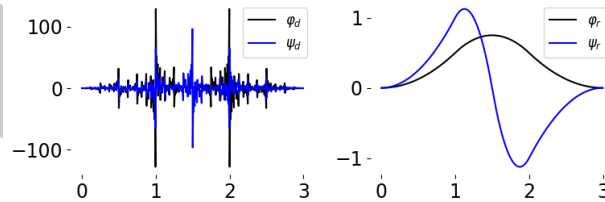


Figure 1: Wavelets and scaling functions of bior3.1 analysis (left) and synthesis (right).

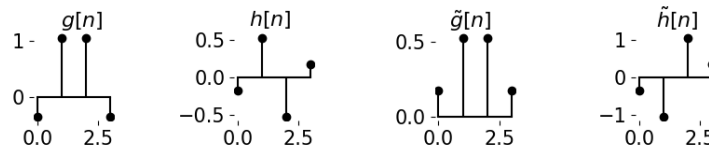


Figure 2: Impulse responses of different bior3.1 analysis (left two) and synthesis (right two) lowpass and highpass FIR filters.

3.1 Circular convolution operators

The four matrices in the two band perfect reconstruction filter bank in Figure 3 are — (i) G is lowpass analysis matrix, (ii) H is highpass analysis matrix, (iii) \tilde{G} is lowpass synthesis matrix and (iv) \tilde{H} is highpass synthesis matrix. These matrices are operators for circular convolution, constructed by circular shifts of the corresponding FIR filters $g[n - k]$, $h[n - k]$, $\tilde{g}[n - k]$ and $\tilde{h}[n - k]$, where $g = \tilde{g}$ and $h = \tilde{h}$ for orthogonal wavelets.

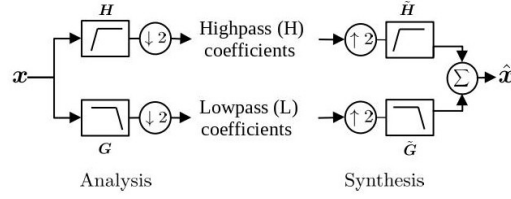


Figure 3: Two band perfect reconstruction filter bank.

3.1.1 Analysis matrix

For a lowpass analysis FIR filter of length L and a input sequence of length N , the circular convolution operator is a matrix G of shape $N \times N$. A downsampling by two $f_{(\downarrow 2)}$ of the output the convolution is equivalent to getting rid of the even rows of G to give $\frac{N}{2} \times N$ operator $f_{(\downarrow 2)}G$. Similarly, for the highpass analysis FIR filter of same wavelet with convolution and downsampling is a $\frac{N}{2} \times N$ operator $f_{(\downarrow 2)}H$. The analysis matrix is,

$$A = \begin{bmatrix} f_{(\downarrow 2)}G \\ f_{(\downarrow 2)}H \end{bmatrix}_{N \times N}$$

where A is a $N \times N$ decimated circular convolution operator formed by combining the lowpass and highpass decimated operators].

3.1.2 Synthesis matrix

The synthesis matrix S is another $N \times N$ decimated circular convolution operator as given by equation [eq:SynthesisMatrix],

$$S = \begin{bmatrix} f_{(\downarrow 2)}\tilde{G} \\ f_{(\downarrow 2)}\tilde{H} \end{bmatrix}_{N \times N}^T$$

where \tilde{G} and \tilde{H} are matrices formed by the lowpass and highpass synthesis FIR filters. The above is a general representation for both orthogonal and biorthogonal wavelets families where for orthogonal wavelets, $\tilde{G} = G$, $\tilde{H} = H$ and thus $S = A^T$.

A two-band perfect reconstruction discrete wavelet system for one-dimensional inputs is given by the analysis equation and the synthesis equation,

$$q = \text{DWT}(x) \text{ or } q = (Ax^T)^T \text{—Analysis}$$

$$x = \text{IDWT}(q) \text{ or } x = (Sq^T)^T \text{—Synthesis}$$

where A and S are analysis and synthesis matrices, x is a input sequence and q has a distinct lowpass and a highpass subband.

Example 1. Given, a sequence $x \in \mathbb{R}^8$ or $N = 8$ and FIR filters length $L = 6$.

$$\text{LPF and downsampling } f_{(\downarrow 2)}G = \begin{bmatrix} g_1 & g_0 & 0 & 0 & g_5 & g_4 & g_3 & g_2 \\ g_3 & g_2 & g_1 & g_0 & 0 & 0 & g_5 & g_4 \\ g_5 & g_4 & g_3 & g_2 & g_1 & g_0 & 0 & 0 \\ 0 & 0 & g_5 & g_4 & g_3 & g_2 & g_1 & g_0 \end{bmatrix}_{\frac{N}{2} \times N}$$

HPF and downsampling $f_{(\downarrow 2)} \mathbf{H} = \begin{bmatrix} h_1 & h_0 & 0 & 0 & h_5 & h_4 & h_3 & h_2 \\ h_3 & h_2 & h_1 & h_0 & 0 & 0 & h_5 & h_4 \\ h_5 & h_4 & h_3 & h_2 & h_1 & h_0 & 0 & 0 \\ 0 & 0 & h_5 & h_4 & h_3 & h_2 & h_1 & h_0 \end{bmatrix}_{\frac{N}{2} \times N}$

Analysis matrix is $\mathbf{A} = \begin{bmatrix} g_1 & g_0 & 0 & 0 & g_5 & g_4 & g_3 & g_2 \\ g_3 & g_2 & g_1 & g_0 & 0 & 0 & g_5 & g_4 \\ g_5 & g_4 & g_3 & g_2 & g_1 & g_0 & 0 & 0 \\ 0 & 0 & g_5 & g_4 & g_3 & g_2 & g_1 & g_0 \\ h_1 & h_0 & 0 & 0 & h_5 & h_4 & h_3 & h_2 \\ h_3 & h_2 & h_1 & h_0 & 0 & 0 & h_5 & h_4 \\ h_5 & h_4 & h_3 & h_2 & h_1 & h_0 & 0 & 0 \\ 0 & 0 & h_5 & h_4 & h_3 & h_2 & h_1 & h_0 \end{bmatrix}_{N \times N}$

92 Similarly,

Synthesis matrix is $\mathbf{S} = \begin{bmatrix} \tilde{g}_1 & \tilde{g}_0 & 0 & 0 & \tilde{g}_5 & \tilde{g}_4 & \tilde{g}_3 & \tilde{g}_2 \\ \tilde{g}_3 & \tilde{g}_2 & \tilde{g}_1 & \tilde{g}_0 & 0 & 0 & \tilde{g}_5 & \tilde{g}_4 \\ \tilde{g}_5 & \tilde{g}_4 & \tilde{g}_3 & \tilde{g}_2 & \tilde{g}_1 & \tilde{g}_0 & 0 & 0 \\ 0 & 0 & \tilde{g}_5 & \tilde{g}_4 & \tilde{g}_3 & \tilde{g}_2 & \tilde{g}_1 & \tilde{g}_0 \\ \tilde{h}_1 & \tilde{h}_0 & 0 & 0 & \tilde{h}_5 & \tilde{h}_4 & \tilde{h}_3 & \tilde{h}_2 \\ \tilde{h}_3 & \tilde{h}_2 & \tilde{h}_1 & \tilde{h}_0 & 0 & 0 & \tilde{h}_5 & \tilde{h}_4 \\ \tilde{h}_5 & \tilde{h}_4 & \tilde{h}_3 & \tilde{h}_2 & \tilde{h}_1 & \tilde{h}_0 & 0 & 0 \\ 0 & 0 & \tilde{h}_5 & \tilde{h}_4 & \tilde{h}_3 & \tilde{h}_2 & \tilde{h}_1 & \tilde{h}_0 \end{bmatrix}_{N \times N}$

93 The DWT of x produces subbands,

$$\mathbf{q} = \text{DWT}(x) \text{ or, } \mathbf{q} = \mathbf{A}x$$

94 Perfect reconstruction,

$$x = \text{IDWT}(\mathbf{q}) \text{ or, } x = \mathbf{S}q$$

95 3.2 DWT 1D layer

96 A DWT 1Dayer operates on input tensors of shape (batch, length, channels) and produces an
97 output of shape (batch, length/2, 2 × channels). as described in Algorithm 1a.

98 **Algorithm 1a** —

- 99 1. Input \mathbf{X} of shape (batch, length, channels).
- 100 2. Generate analysis matrix \mathbf{A} using length of input.
- 101 3. For each batched channel $x \in \mathbf{X}$ of shape (batch, length): $\mathbf{q}_c = \mathbf{A}x^T$
- 102
- 103 4. Stacking for all c channels: $\mathbf{Q} := (\mathbf{q}_c)_{\vee c}$ to a shape (batch, length, channels).
- 104 5. Group subbands and return an output $\mathbf{Q}^{(\text{grouped})}$ of shape (batch, length/2, 2 × channels)

Grouping two subbands in DWT 1D

mid = int(Q.shape[1]/2)

L = Q[:, :mid, :]

H = Q[:, mid:, :]

out = Concatenate([L, H], axis=-1)

3.3 IDWT 1D layer

An IDWT 1D layer operates on input tensors of shape (batch, length/2, $2 \times$ channels) and produces an output of shape (batch, length, channels) as described in Algorithm 1b.

Algorithm 1b —

1. Input $Q^{(\text{grouped})}$ of shape (batch, length/2, $2 \times$ channels)
2. Ungroup the subbands to get Q of shape (batch, length, channels)
3. Generate synthesis matrix S using length of Q .
4. For each batched channel $q \in Q$ of shape (batch, length): $x = Sq^T$, i.e., a perfect reconstruction, where, $S = A^T$ for orthogonal wavelets
5. Layer output (perfect reconstruction): $X := (x_c)_{\forall c}$ is of shape (batch, length, channels)

4 Higher dimensional discrete wavelet systems

In sequences (1D), the DWT 1D applies along the only independent variable. To achieve higher dimensional DWT, the same DWT 1D is successively applied separately to all the independent variables. For example, in an image (2D) the DWT 2D is a row-wise DWT 1D followed by a column-wise DWT 1D. Similarly, the reconstruction is column-wise IDWT 1D followed by a row-wise IDWT 1D.

4.1 Two-dimensional discrete wavelet system

The pixel values in an image are a function of two independent spatial axes. A DWT 2D filter bank is a separable transform with row-wise filtering followed by column-wise filtering that yields four subbands - LL, LH, HL and HH. A two-dimensional discrete wavelet system is,

$$q = \text{DWT}(x) := A(Ax_{021})_{021}^T \quad \text{—Analysis}$$

$$x = \text{IDWT}(q) := S(S[A(Ax_{021})_{021}^T]_{021}^T)_{021}^T \quad \text{—Synthesis}$$

where, A and S are the same analysis and synthesis matrices defined for one-dimensional wavelet system. Figure 4 shows an $N \times N$ image and its four localized spatial-frequency subbands after DWT. Here, the low-frequency band is LL, and the other three are high-frequency subbands representing horizontal, vertical and diagonal features. Figure 5 illustrates a separable 2D DWT perfect reconstruction filter bank. The 2D layers operate on batched, multichannel tensors of shape (batch, height, width, channels), where each image is of shape height and width. Figure 6 illustrates input, output and perfect reconstruction by DWT 2D and IDWT 2D layers. The Multiresolution Encoder-Decoder Convolutional Neural Network in (Tarafdar et al., 2025) uses these forward and inverse transform layers.

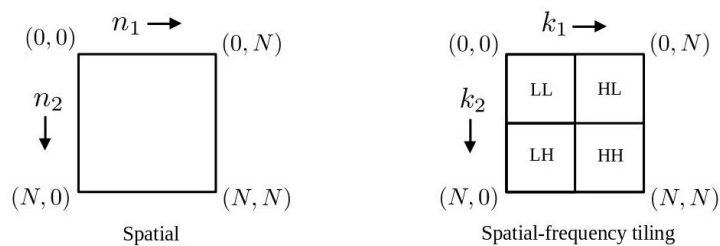


Figure 4: Natural domain (left) and spatial-frequency tiling (right) after a DWT 2D.

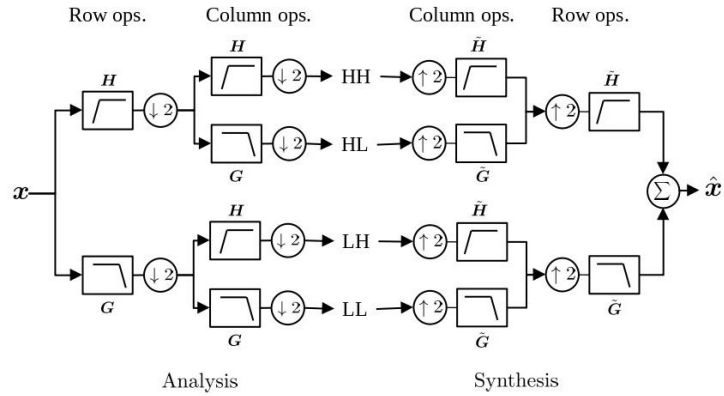


Figure 5: Separable DWT 2D perfect reconstruction filter bank.

4.1.1 DWT 2D layer

A DWT 2D layer operates on input tensors of shape (batch, height, width, channels) and produces an output of shape (batch, height/2, width/2, 4 × channels) as described in Algorithm 2a.

Algorithm 2a —

1. Input X of shape (batch, height, width, channels).
2. Generate analysis matrix A using height and width of input.
3. For each batched channel $x_c \in X$ of shape (batch, height, width):
(omitting suffix c in x below for simplicity of notation)
 - a. Row-wise batch DWT 1D: $Ax_{021}^T := \text{Einsum}(ij, bjk \rightarrow bik)$
 - b. Column-wise batch DWT 1D: $A(Ax_{021}^T)_{021}^T := \text{Einsum}(ij, bjk \rightarrow bik)$ Or, equivalently, DWT of a batched channel x is,

$$q_c = \text{DWT}(x) \text{ or } q_c = A(Ax_{021})_{021}^T$$

where, the suffix 021 in x_{021}^T denotes permutation of axis, i.e., transpose.

4. Stacking for all c channels: $Q := (q_c)_{\forall c}$ to a shape (batch, height, width, channels).
5. Group subbands and return an output $Q^{(\text{grouped})}$ of shape (batch, height/2, width/2, 4 × channels)

Grouping four subbands in DWT 2D

```
mid = int(Q.shape[1]/2)
LL = Q[:, :mid, :mid, :]
LH = Q[:, mid:, :mid, :]
HL = Q[:, :mid, mid:, :]
HH = Q[:, mid:, mid:, :]
output = Concatenate([LL, LH, HL, HH], axis=-1)
```

4.1.2 IDWT 2D layer

An IDWT 2D layer operates on input tensors of shape (batch, height/2, width/2, 4 × channels) and produces an output of shape (batch, height, width, channels) as described in Algorithm 2b.

Algorithm 2b —

1. Input $Q^{(\text{grouped})}$ of shape (batch, height/2, width/2, 4 × channels)

2. Ungroup the subbands to get Q of shape (batch, height, width, channels)
3. Generate synthesis matrix S using height and width of Q .
4. For each batched channel $q_c \in Q$ of shape (batch, height, width):
(omitting suffix c in q below for simplicity of notation)
 - a. Row-wise batch IDWT 1D: $Sq_{021}^T := \text{Einsum}(ij, bjk \rightarrow bik)$
 - b. Column wise batch IDWT 1D: $S(Sq_{021}^T)_{021}^T := \text{Einsum}(ij, bjk \rightarrow bik)$ or equivalently, a perfect reconstruction,

$$x = \text{IDWT}(q) \text{ or } x = S(Sq_{021}^T)_{021}^T$$

where, the suffix 021 in x_{021}^T denotes permutation of axis, i.e., transpose and $S = A^T$ for orthogonal wavelets

5. Layer output : $X := (x_c)_{\forall c}$ is of shape (batch, height, width, channels)
(Perfect reconstruction)

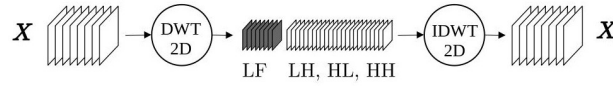


Figure 6: DWT decomposition and perfect reconstruction of a multichannel image tensor.

4.2 Three-dimensional discrete wavelet system

A three-dimensional (3D) discrete wavelet system for a 3D input x is given by,

$$q = \text{DWT}(x) := [A(A(Ax_{0213}^T)_{0213}^T)_{0132}^T]_{0132}^T \quad \text{—Analysis}$$

$$x = \text{IDWT}(q) := (S(S(S[A(A(Ax_{0213}^T)_{0213}^T)_{0132}^T]_{0132}^T)_{0132}^T)_{0213}^T)_{0213} \quad \text{—Synthesis}$$

where, A and S are the same analysis and synthesis matrices as defined for one-dimensional wavelet system. The DWT 3D and IDWT 3D layers operate on batched, multichannel tensors of shape (batch, height, width, depth, channels).

4.2.1 DWT 3D layer

A DWT 3D layer operates on input tensors of shape (batch, height, width, depth, channels) and produces an output of shape (batch, height/2, width/2, depth/2, $8 \times$ channels) as described in Algorithm 3a.

Algorithm 3a —

1. Input X of shape (batch, height, width, depth, channels).
2. Generate analysis matrix A using height, width and depth of input.
3. For each batched channel $x_c \in X$ of shape (batch, height, width, depth):
(omitting suffix c in x below for simplicity of notation)
 - a. Row-wise batch DWT 1D: $Ax_{0213}^T := \text{Einsum}(ij, bjk \rightarrow bik)$
 - b. Column-wise batch DWT 1D: $A(Ax_{0213}^T)_{0213}^T := \text{Einsum}(ij, bjk \rightarrow bik)$
 - c. Depth-wise batch IDWT 1D: $A(A(Ax_{0213}^T)_{0213}^T)_{0132}^T := \text{Einsum}(ik, bjk \rightarrow bjil)$
 Therefore, DWT of x yield coefficients:

$$q_c := \text{DWT}(x) = [A(A(Ax_{0213}^T)_{0213}^T)_{0132}^T]_{0132}^T$$

- 185 4. Stacking for all c channels as $Q := (q_c)_{\forall c}$ to a shape (batch, height, width, depth,
186 channels).
- 187 5. Group subbands and return an output $Q^{(\text{grouped})}$ of shape (batch, height/2, width/2,
188 depth/2, $8 \times \text{channels}$).

```
# Grouping Eight subbands in 3D DWT
mid = int(Q.shape[2]/2)
LLL = Q[:, :mid, :mid, :mid, :]
LLH = Q[:, :mid, :, :mid, :mid, :]
LHL = Q[:, :mid, mid, :, :mid, :]
LHH = Q[:, :mid, mid, :, :mid, :]
HLL = Q[:, :mid, :mid, mid, :, :]
HLH = Q[:, :mid, :mid, mid, :, :]
HHL = Q[:, :mid, mid, mid, :, :]
HHH = Q[:, :mid, mid, mid, :, :]
output = Concatenate([LLL, LLH, LHL, LHH, HLL, HLH, HHL, HHH], axis=-1)
```

189 4.2.2 IDWT 3D layer

190 An IDWT 3D layer operates on input tensors of shape (batch, height/2, width/2, depth/2, $8 \times$
191 channels) and produces an output of shape (batch, height, width, depth, channels) as de-
192 scribed in Algorithm 3b.

193 Algorithm 3b —

- 194 1. Input $Q^{(\text{grouped})}$ of shape (batch, height/2, width/2, depth/2, $8 \times \text{channels}$)
- 195 2. Ungroup to get Q of shape (batch, height, width, depth, channels)
- 196 3. Generate synthesis matrix S using height, width and depth of Q .
- 197 4. For each batched channel $q_c \in Q$ of shape (batch, height, width, depth):
198 (omitting suffix c in q below for simplicity of notation)
- 199 a. Row-wise batch IDWT 1D: $Sq_{0132}^T := \text{Einsum}(ik, b j k l \rightarrow b j i l)$
- 200 b. Column-wise batch IDWT 1D: $S(Sq_{0132}^T)_{0132} := \text{Einsum}(ij, b j k l \rightarrow b i k l)$
- 201 c. Depth-wise batch IDWT 1D: $S(S(Sq_{0132}^T)_{0132}^T)_{0213} := \text{Einsum}(ij, b j k l \rightarrow b i k l)$
202 or equivalently, a perfect reconstruction,

$$x_c := \text{IDWT}(q) = [S(S(Sq_{0132}^T)_{0132}^T)_{0213}]_{0213}^T$$

204 where, $S = A^T$ for orthogonal wavelets.

- 205 5. Layer output : $X := (x_c)_{\forall c}$ is of shape (batch, height, width, depth, channels)
206 (Perfect reconstruction)

207 In general, a seamless realization of fast D-dimensional DWT and IDWT is possible by extending
208 the above separable method to all the independent N axes one after the other. The number
209 of subbands will be equal to 2^D for a D dimensional DWT. For example, sequences ($D = 1$)
210 yield two subbands, images ($D = 2$) yield four subbands, three-dimensional inputs ($D = 3$)
211 with voxels yield eight subbands etc.

212 5 Multilevel wavelet filter banks

213 The above-discussed DWT and IDWT layers are building blocks in constructing multilevel
214 DWT filter banks. Figure 7 shows the partitioning of the 1D frequency axis and tiling of
215 the 2D frequency plane using a level-4 1D and 2D DWT. The multilevel DWT successively
216 decomposes the low-frequency feature. If the high-frequency features are also decomposed
217 successively, then we get a Wavelet Packet Transform (WPT) filter bank.

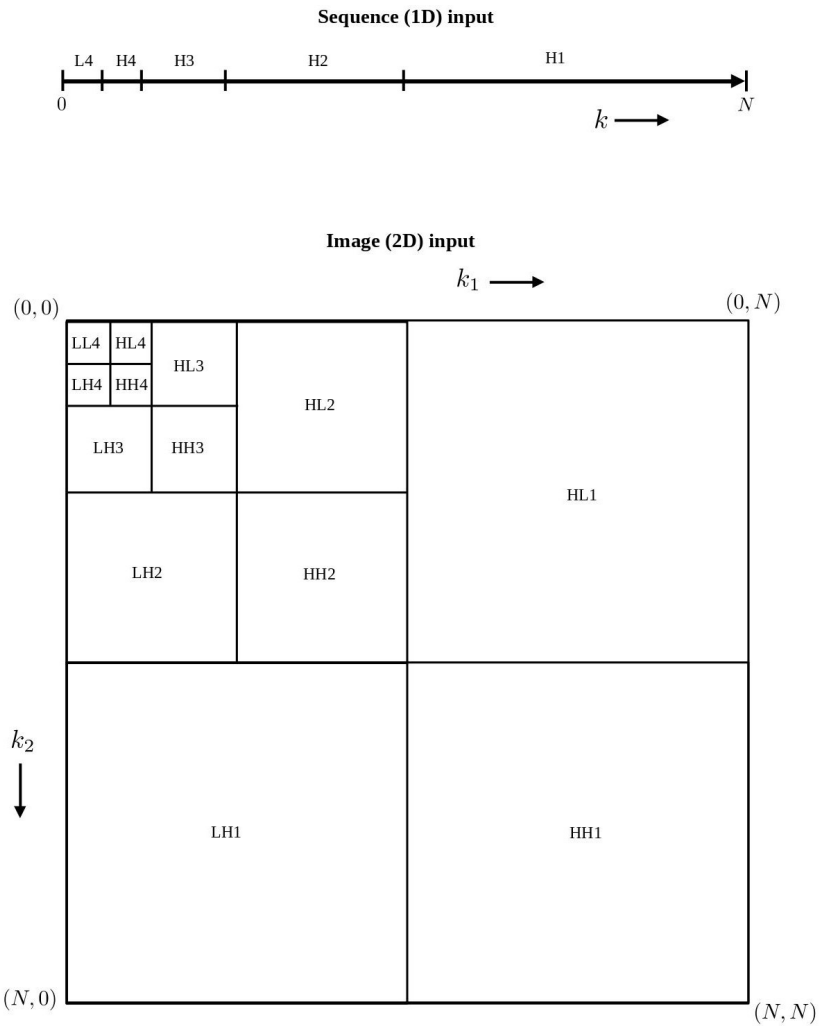


Figure 7: Spatial-frequency tiling by DWT level 4 decomposition of a sequence of length N (top) and image of shape Nimes N (bottom).

References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., ... Zheng, X. (2015). *TensorFlow: Large-scale machine learning on heterogeneous systems*. <https://www.tensorflow.org/>
- Cotter, F. (2023). *Pytorch-wavelets: A port of the DTCWT toolbox to run on pytorch*. <https://pypi.org/project/pytorch-wavelets/>
- Fatica, M. (2008). CUDA toolkit and libraries. *2008 IEEE Hot Chips 20 Symposium (HCS)*, 1–22. <https://doi.org/10.1109/HOTCHIPS.2008.7476520>
- Harris, C. R., Millman, K. J., Walt, S. J. van der, Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., Kerkwijk, M. H. van, Brett, M., Haldane, A., Río, J. F. del, Wiebe, M., Peterson, P., ... Oliphant, T. E. (2020). Array programming with NumPy. *Nature*, 585(7825), 357–362. <https://doi.org/10.1038/s41586-020-2649-2>
- Imambi, S., Prakash, K. B., & Kanagachidambaresan, G. (2021). PyTorch. *Programming*

- 233 with TensorFlow: Solution for Edge Computing Applications, 87–104. [https://doi.org/10.](https://doi.org/10.1007/978-3-030-57077-4_10)
234 1007/978-3-030-57077-4_10
- 235 Lee, G., Gommers, R., Waselewski, F., Wohlfahrt, K., & O’Leary, A. (2019). PyWavelets:
236 A python package for wavelet analysis. *Journal of Open Source Software*, 4(36), 1237.
237 <https://doi.org/10.21105/joss.01237>
- 238 Leiderman, T. (2025). *Tensorflow-wavelets: Tensorflow wavelet layers*. [https://pypi.org/](https://pypi.org/project/tensorflow-wavelets/)
239 [project/tensorflow-wavelets/](https://pypi.org/project/tensorflow-wavelets/)
- 240 Misiti, M., Misiti, Y., Oppenheim, G., & Poggi, J.-M. (1996). Wavelet toolbox. *The*
241 *MathWorks Inc., Natick, MA*, 15, 21.
- 242 Tarafdar, K. K., Meher, A., Shah, M., Saif, Q., Kumar, D., Nimkar, A. V., Jayasundar, R.,
243 & Gadre, V. M. (2025). Multiresolution encoder-decoder convolutional neural network for
244 magnetic resonance image segmentation. *ICASSP 2025-2025 IEEE International Conference*
245 *on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. [https://doi.org/10.1109/](https://doi.org/10.1109/ICASSP49660.2025.10890832)
246 [ICASSP49660.2025.10890832](https://doi.org/10.1109/ICASSP49660.2025.10890832)
- 247 Versaci, F. (2021). WaveTF: A fast 2D wavelet transform for machine learning in keras. *Pattern*
248 *Recognition. ICPR International Workshops and Challenges: Virtual Event, January 10–15,*
249 *2021, Proceedings, Part i*, 605–618. https://doi.org/10.1007/978-3-030-68763-2_46

DRAFT