# Spectral Wavelet Dropout: Regularization in the Wavelet Domain

Rinor Cakaj Signal Processing Robert Bosch GmbH & University of Stuttgart 71229 Leonberg, Germany Rinor.Cakaj@de.bosch.com Jens Mehnert Signal Processing Robert Bosch GmbH 71229 Leonberg, Germany JensEricMarkus.Mehnert@de.bosch.com Bin Yang ISS University of Stuttgart 70550 Stuttgart, Germany bin.yang@iss.uni-stuttgart.de

Abstract—Regularization techniques help prevent overfitting and therefore improve the ability of convolutional neural networks (CNNs) to generalize. One reason for overfitting is the complex co-adaptations among different parts of the network, which make the CNN dependent on their joint response rather than encouraging each part to learn a useful feature representation independently. Frequency domain manipulation is a powerful strategy for modifying data that has temporal and spatial coherence by utilizing frequency decomposition. This work introduces Spectral Wavelet Dropout (SWD), a novel regularization method that includes two variants: 1D-SWD and 2D-SWD. These variants improve CNN generalization by randomly dropping detailed frequency bands in the discrete wavelet decomposition of feature maps. Our approach distinguishes itself from the preexisting Spectral "Fourier" Dropout (2D-SFD), which eliminates coefficients in the Fourier domain. Notably, SWD requires only a single hyperparameter, unlike the two required by SFD. We also extend the literature by implementing a one-dimensional version of Spectral "Fourier" Dropout (1D-SFD), setting the stage for a comprehensive comparison. Our evaluation shows that both 1D and 2D SWD variants have competitive performance on CIFAR-10/100 benchmarks relative to both 1D-SFD and 2D-SFD. Specifically, 1D-SWD has a significantly lower computational complexity compared to 1D/2D-SFD. In the Pascal VOC Object Detection benchmark, SWD variants surpass 1D-SFD and 2D-SFD in performance and demonstrate lower computational complexity during training.

*Index Terms*—Regularization, Overfitting, Generalization, Wavelets, Fourier, Dropout

## I. INTRODUCTION

Over recent years, convolutional neural networks (CNNs) have demonstrated significant advancements in a variety of computer vision applications, such as image recognition [1] or object detection [2]. However, CNNs are prone to overfitting, especially when the available training data is limited or when the model's capacity is excessively large [3]. Regularization techniques are necessary to ensure that models can generalize effectively to new, unseen data [4].

Various regularization methods have been developed to address the overfitting problem in CNNs. These include weight penalties such as  $L_1$  regularization [5] and weight decay [6], in addition to techniques like data augmentation [7], and ensemble learning methods [8]. Among these, dropout mechanisms, notably standard dropout [9] and dropconnect [3], play an important role by randomly dropping neurons, weights, or feature map entries during training to prevent co-adaptation of units within the network. However, the effectiveness of these methods is often compromised in convolutional layers where features have spatial correlation, leading to suboptimal regularization and the continuation of overfitting problems [10].

The limitations within standard dropout approaches have led to the development of more advanced variants like dropblock [10] and dropout2d [11]. These methods introduce structured forms of dropout, designed to improve regularization in CNNs by strategically dropping feature map segments or channels. This addresses the spatial correlation in the feature maps and therefore increases the regularization effect.

The investigation of alternative regularization strategies has led to the use of frequency domain manipulation. Frequency decomposition provides a natural framework for the manipulation of data with spatial coherence [12]. Khan et al. [13] introduced Spectral "Fourier" Dropout (2D-SFD), a strategy designed to reduce overfitting by selectively pruning and dropping Fourier domain coefficients within feature maps. The process begins with the application of a two-dimensional discrete cosine transform (2D-DCT) to transform each channel in feature maps into the frequency domain. It then prunes Fourier coefficients that are considered weak or "noisy", i.e. those that fall below the  $\eta$  quantile. Subsequently, a dropout mask is applied to the remaining coefficients with a dropout probability p. Finally, the modified feature map is transformed back into the spatial domain using an inverse 2D-DCT. The effectiveness of 2D-SFD lies in its ability to make CNNs insensitive to weak and "noisy" spectral components. This is achieved by randomly preserving only the important Fourier coefficients for signal reconstruction. Moreover, it effectively prevents co-adaptation among feature detectors by diversifying the network's reliance across various frequency components [13].

The discrete wavelet transform (DWT) provides an alternative way for regularization, offering advantages in handling multi-resolution data. Unlike the Fourier transform, which decomposes signals into sinusoidal components without localization, the DWT divides data into frequency bands at multiple resolutions, capturing both global and localized features [14]. This capability motivates the investigation of Spectral Wavelet Dropout (SWD), a novel method proposed in this work for

© 2024 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.



Fig. 1. Illustration of the 2D Spectral Wavelet Dropout process applied to an input feature map. A one-level 2D wavelet decomposition divides the map into four sub-bands: LL (low-frequency), LH (horizontal high-frequency), HL (vertical high-frequency), and HH (diagonal high-frequency), presented sequentially from top to bottom in the figure. The high-frequency bands are randomly dropped (HL is shown as dropped) according to a dropout probability p, while the LL band is consistently maintained. The feature map is then transformed back into the spatial domain using an inverse 2D wavelet transform (2D-IDWT).

improving CNN generalization by randomly dropping wavelet bands within the discrete wavelet decomposition of feature maps.

We introduce two variants of Spectral Wavelet Dropout: the one-dimensional (1D-SWD) and two-dimensional (2D-SWD) versions. For 1D-SWD, consider a feature map  $X \in \mathbb{R}^{C \times H \times W}$ . The initial step involves flattening the feature map across each channel, resulting in  $\hat{X} \in \mathbb{R}^{C \times H \cdot W}$ . Then, 1D-SWD performs a three-level one-dimensional discrete wavelet decomposition (1D-DWT) of  $\hat{X}$ . This decomposition produces approximation and detailed frequency coefficients across three levels, also referred to as frequency bands. In the next phase, 1D-SWD randomly drops detailed wavelet bands based on a dropout probability p, effectively introducing regularization. The final step involves transforming the feature map back into the spatial domain through an inverse 1D-DWT and reshaping it for input into the subsequent network layer. An illustration of 1D-SWD is shown in Figure 3 in the Appendix.

The 2D-SWD method applies a one-level two-dimensional wavelet decomposition to each channel of the feature map X. This decomposition process produces four distinct sub-bands: LL (approximation), LH (horizontal details), HL (vertical details), and HH (diagonal details). Following the 1D-SWD method, 2D-SWD randomly drops the LH, HL, and HH detailed frequency bands, based on a dropout probability p, while preserving the LL approximation band. The modified feature map is then transformed back into the spatial domain through an inverse 2D wavelet transform (2D-IDWT) before being input into the subsequent network layer. Both 1D-SWD and 2D-SWD are characterized by a single hyperparameter,

p, which considerably simplifies the hyperparameter tuning process. An illustration of 2D-SWD is shown in Figure 1.

Jo and Bengio [15] showed that CNNs have a bias towards higher frequency components. Similarly, Geirhos et al. [16] presented a texture bias in CNNs trained on ImageNet [17]. Building on these findings, our method differs from the Spectral "Fourier" Dropout [13] approach by strategically targeting the frequency domain. Specifically, SWD randomly drops detailed frequency bands. This approach is based on the understanding that the approximation band of the wavelet transformation typically provides a smoothed representation of the original signal, while the detailed bands capture high frequency variation and noise. By randomly dropping the detailed bands, SWD encourages the CNN to diversify its reliance away from single detailed frequency bands within the feature maps.

We extend the literature by implementing a one-dimensional version of Spectral "Fourier" Dropout (1D-SFD), to present a fair comparison of the methods. As a one-dimensional analogue to 2D-SFD, 1D-SFD begins with the flattening of the feature map across each channel. This is followed by the application of a one-dimensional discrete cosine transform (1D-DCT). Subsequently, a specified proportion of activations, denoted by  $\eta$ , is pruned, and a dropout mask is applied to the remaining Fourier coefficients using a dropout probability *p*. Finally, the modified feature map is transformed back to the spatial domain using the inverse 1D-DCT. It is noteworthy that 1D-SFD uses two hyperparameters.

The motivation behind using 1D transformations lies in their computational efficiency. Moreover, they offer specific pattern

recognition capabilities for images with informative horizontal features. While 2D transformations provide a more detailed analysis by capturing both vertical and horizontal patterns, our experiments demonstrate that 1D approaches can achieve comparable regularization effects with lower computational demands.

Our evaluation shows that both 1D and 2D-SWD variants have competitive performance on CIFAR-10/100 benchmarks relative to both 1D-SFD and 2D-SFD. Specifically, 1D-SWD is highlighted for its competitive results, considering its significantly lower computational complexity during training compared to 1D/2D-SFD. On ImageNet, 1D-SWD maintains competitive performance considering the reduced computational demand during training. In the Pascal VOC Object Detection benchmark, SWD variants exceed 1D-SFD and 2D-SFD in terms of performance while having a lower training overhead. Specifically, 2D-SWD improves the mean Average Precision (mAP) by 0.56 percentage points (a 0.72% increase) compared to the baseline. A key benefit of the SWD method is its simplified setup, which only requires one hyperparameter, unlike the SFD method that needs two.

## II. RELATED WORK

#### A. Regularization

Regularization is one of the key elements of deep learning [4], allowing the model to generalize well to unseen data even when trained on a finite training set or with an imperfect optimization procedure [18]. There are several techniques to regularize CNNs which can be categorized into groups. Data augmentation methods like cropping, flipping, adjusting brightness or sharpness [7] and cutout [19] transform the training dataset to avoid overfitting. Regularization techniques like dropout [9], dropblock [10] or dropconnect [3] drop neurons or weights from the CNN during training to prevent units from co-adapting too much [9]. Khan et al. [13] extend the idea of dropping parts of CNNs to the spectral domain. Spectral "Fourier" Dropout avoids overfitting by pruning weak and "noisy" Fourier domain coefficients of the feature maps and randomly dropping a percentage of the remaining Fourier domain coefficients.

Furthermore, CNNs can be regularized using penalty terms in the loss function. Weight decay [6] encourages the weights of the CNN to be small in magnitude. The  $L_1$ -regularization [5] encourages the weights of non-relevant features to be zero.

## B. Wavelet Transforms in CNNs

Wavelet transformations are used in image processing [20], for example the JPEG2000 format [21] uses wavelets to represent images as highly sparse feature maps. Recently, wavelet transforms have been used in deep learning due to their multi-resolution analysis properties.

1) Wavelet Neural Networks: In a limited form, CNNs can be seen as a form of multi-resolution analysis. However, conventional CNNs are missing a large part of spectral information present in feature maps [22]. Therefore, different

approaches supplement the missing part of the spectral information via additional components in the entire architecture using the DWT. Fujieda et al. [22] present wavelet CNNs and demonstrate benefits in texture classification. Liu et al. [23] integrates a wavelet transform into the CNN architecture to reduce the resolution of feature maps, while at the same time increasing the receptive field.

2) *Pooling:* Depending on the data, max pooling can remove and average pooling can dilute important details from an image. Wavelet Pooling [24] overcomes these issues by decomposing the feature maps into a second-level decomposition and removing the first-level sub-bands to reduce the feature dimensions. Adaptive wavelet pooling [25] extends this approach by using adaptive- and scaled-wavelets in the pooling process.

3) Compression: CNNs are successfully used in a number of applications. Despite this, their storage requirements have largely prevented their use on mobile devices. Wolter et al. [26] showed how the fast wavelet transform can be used to compress linear layers in neural networks using learnable wavelets.

## III. METHOD

One of the main advantages of the wavelet transform over the Fourier transform is its superior spatial/temporal resolution. This property allows the wavelet transform to capture information about both frequency and location simultaneously, providing an optimal structure for manipulating data with spatial coherence.

Inspired by the concept of Spectral "Fourier" Dropout presented by Khan et al. [13], we introduce Spectral Wavelet Dropout, which includes two variants: 1D-SWD and 2D-SWD. These variants improve CNN generalization by randomly dropping detailed frequency bands in the discrete wavelet decomposition of feature maps. To effectively integrate SWD into CNNs, we first look at the basics of one- and twodimensional discrete wavelet transforms.

## A. Discrete Wavelet Transform

The discrete wavelet transform is a method for analyzing signals and images using discretely sampled wavelets. The DWT can be implemented with various wavelet families, such as Haar wavelets [27] and Daubechies wavelets [28].

Given a vector  $x \in \mathbb{R}^d$  of even length, the one-level one-dimensional DWT (1D-DWT) is computed by processing x through corresponding low-pass and high-pass filters, denoted as g and h. These filters, which are quadrature mirror filters [29], effectively split the frequency spectrum of the signal, allowing every other sample to be removed according to Nyquist's theorem [30]. The computation is formalized by the following equations:

$$y_{\text{low}}[n] = (g * x)[n] = \sum_{\substack{k=0\\N-1}}^{N-1} g[k] \cdot x[2n-k], \qquad (1)$$

$$y_{\text{high}}[n] = (h * x)[n] = \sum_{k=0}^{N-1} h[k] \cdot x[2n-k],$$
 (2)



Fig. 2. Detailed schematic of a three-level wavelet filter bank used in our proposed 1D-SWD method. The filter bank decomposes an input signal x[n] into hierarchical sets of approximation and detailed coefficients. This allows analysis of the signal across multiple resolutions. This structured arrangement, using sequential high-pass h[n] and low-pass g[n] filters with subsequent downsampling, enables SWD to prevent overfitting. By randomly dropping high-frequency details, it promotes model generalization and prevents overfitting.

where N is the filter length and n denotes the output terms, with  $n \in \{1, \ldots, d/2\}$ . The output  $y_{\text{low}}$  represents the approximation coefficients from the low-pass filter and  $y_{\text{high}}$  the detailed coefficients from the high-pass filter.

1D-DWT requires signals with lengths that are powers of two. To address border distortions and signal lengths that do not meet this condition, the signal is typically extended. One common method is zero-padding, where zeros are appended to the signal to meet the required length for transformation.

The 1D-DWT's decomposition can be iteratively applied to the approximation coefficients to refine frequency resolution. Each iteration applies high- and low-pass filters to the signal followed by downsampling. This process can be visualized in a binary tree structure, often referred to as a filter bank (as shown in Figure 2). The downsampling operation is symbolized by  $\downarrow$  2. The approximation coefficients are a smoothed version of the original signal, whereas the detailed coefficients at levels 1, 2, and 3 capture increasingly fine-grained high-frequency content and noise.

The two-dimensional discrete wavelet transform (2D-DWT) is used to analyze signals in two-dimensional spaces, such as images. This transformation is applied sequentially along two orthogonal dimensions, commonly rows and columns. Firstly, a one-level one-dimensional DWT is applied to each row of the image matrix, decomposing it into low and high-frequency components. Then, this decomposition is applied to each column of the resulting sub-images, leading to four distinct sub-bands for each level of decomposition:

- *LL* (Low-Low) sub-band: This contains the approximation coefficients, resulting from applying the low-pass filter across both dimensions. It shows the primary features of the image, serving as a blurred, downsampled version of the original.
- *LH* (Low-High) sub-band: This contains the vertical details of the image, derived from applying the low-pass filter horizontally and the high-pass filter vertically.
- *HL* (High-Low) sub-band: This contains horizontal details, derived by applying the high-pass filter horizontally

and the low-pass filter vertically.

• *HH* (High-High) sub-band: This sub-band contains diagonal details, derived by applying the high-pass filter across both dimensions.

The DWT operates through matrix multiplication, characterizing it as a linear transformation. This linearity is essential because, during backpropagation, the gradient computation with respect to the DWT is directly determined by the transform matrix itself.

#### B. Spectral Wavelet Dropout

Spectral Wavelet Dropout is introduced in two variants, 1D-SWD and 2D-SWD. Given a batch of feature maps  $X \in \mathbb{R}^{B \times C \times H \times W}$ , 1D-SWD begins by flattening X along the spatial dimensions, resulting in  $\hat{X} \in \mathbb{R}^{B \times C \times H \cdot W}$ . A threelevel 1D-DWT is performed using the Daubechies-3 (db3) wavelet [28], producing approximation and detailed frequency bands.

Subsequently, a mask  $m \in \{0,1\}^3$  is generated, where each element  $m_i$  is drawn from a Bernoulli distribution with a probability of (1-p) for being 1. This mask is used to selectively drop detailed frequency bands. The approximation coefficients are preserved to maintain the primary signal structure. To compensate for the effect of dropout on signal energy, the feature maps are scaled by a factor of  $(1-p)^{-1}$ .

The process for 2D-SWD is analogous. A one-level 2D-DWT is applied to each channel, resulting in LL, LH, HL, and HH sub-bands. The mask is then applied to the detailed sub-bands.

Both processes end with the inverse DWT, 1D-IDWT for 1D-SWD and 2D-IDWT for 2D-SWD, followed by reshaping to the original dimensions in the 1D-DWT variant. Algorithm 1 shows the steps for the 1D-SWD. The complete procedure for the 2D-SWD is explained in Section VII-B the Appendix.

The principle of SWD is based on the ability of the wavelet transformation to decompose signals into components of different spatial-frequency localization. By selectively dropping the detailed frequency components, SWD encourages the CNN to not overly rely on any specific set of frequencies, thereby preventing co-adaptation of feature detectors. This promotes a learning of features across all spatial frequencies present in the data.

#### **IV. EXPERIMENTS**

We evaluated the Spectral Wavelet Dropout method on supervised image recognition on CIFAR-10/100 [31] and ImageNet [17] datasets, and object detection on the Pascal VOC [32] dataset.

The implementation of 1D/2D-SFD [13] was supported by the PyTorch DCT library.<sup>1</sup> For 1D/2D-SWD, we utilized the wavelet transform provided by Cotter [33].

A notable difference in our approach compared to SFD is that 1D/2D-SWD requires only a single hyperparameter - the dropout rate p. We performed a hyperparameter search

<sup>&</sup>lt;sup>1</sup>https://github.com/zh217/torch-dct

## Algorithm 1: Spectral Wavelet Dropout - 1D-SWD

Parameter: dropout probability =  $p \in (0, 1)$ Input:  $X \in \mathbb{R}^{B \times C \times H \times W}$ Output:  $\hat{X} \in \mathbb{R}^{B \times C \times H \times W}$ // Save size of X and flatten the last two dimensions of X 1 size = X.size()

- 2  $\hat{X} = \text{flatten}(X, \text{start}_\text{dim} = 2)$ // Compute the 3-level 1D-DWT for each channel using the daubechie-3 wavelet
- 3 AP, L1, L2, L3 = 1D-DWT $(\hat{X}, J = 3, wave='db3')$ // Create mask for dropout
- 4 dropout\_mask = Bernoulli(3, prob = 1 − p)
  // Wavelet Dropout
- 5  $L1 = L1 \cdot \text{dropout}_{\max}[0]/(1-p)$
- 6  $L2 = L2 \cdot dropout_mask[1]/(1-p)$
- 7  $L3 = L3 \cdot \text{dropout}_{\text{mask}}[2]/(1-p)$
- // Compute inverse Wavelet
   transformation
- 8  $\dot{X} = 1$ D-IDWT([AP, L1, L2, L3], wave='db3') // Reshape the feature map
- 9  $\hat{X} = \operatorname{reshape}(\hat{X}, \operatorname{size})$

for p within the set  $p = \{0.1, 0.2, 0.3, 0.4, 0.5\}$ . In contrast, for 1D/2D-SFD, a grid search was performed over both the dropout rate  $p = \{0.1, 0.2, 0.3, 0.4, 0.5\}$  and the pruning rate  $\eta = \{0, 0.1, 0.2, 0.3, 0.4\}$ . Including  $\eta = 0$  allowed us to evaluate the efficacy of 1D/2D-SFD under a single hyperparameter setting, although this did not consistently yield optimal results. Detailed outcomes of these hyperparameter searches and further implementation details are provided in the Appendix VII-H.

Training overheads are reported relative to the baseline model to offer a clear comparison of the additional computational resources required. Further details on the methodology for runtime measurement are provided in Section V-C.

## A. Image Classification on CIFAR-10/100

We analyzed the impact of both 1D-SWD and 2D-SWD on different architectures, including ResNet18 (R18), ResNet34 (R34), ResNet50 (R50) [1], and VGG16-BN (V16) [34], focusing on accuracy improvements and training overhead. The results are shown in Tables I and II.

Residual Networks (ResNets) [1] are structured into four stages consisting of basic and bottleneck blocks, starting with an initial convolutional layer and ending with a terminal linear layer. Analogous to other dropout variants like dropout2d [11], dropblock [10], and Spectral "Fourier" Dropout [13], we integrate 1D/2D-SWD within the deeper layers of the CNNs to improve regularization. Implementation details are provided in the Appendix, Section VII-D.

1) Results: The results of our experiments are shown in Table I and II. The use of 1D-SWD resulted in accuracy

TABLE I 1D Spectral Dropout Methods on CIFAR-10/100: Accuracy and Training Time Multiplier (TTM).

Method	CIFAR-10 Acc.	CIFAR-100 Acc.	TTM
R18 R18 + 1D-SFD R18 + 1D-SWD	$\begin{array}{c} 94.04\% \pm 0.08\% \\ 94.19\% \pm 0.31\% \\ \mathbf{94.35\% \pm 0.13\%} \end{array}$	$\begin{array}{c} 76.47\% \pm 0.20\% \\ \textbf{77.41\%} \pm \textbf{0.13\%} \\ 77.07\% \pm 0.18\% \end{array}$	1.00x 1.15x 1.12x
R34 R34 + 1D-SFD R34 + 1D-SWD	$\begin{array}{c} 93.69\% \pm 0.30\% \\ 93.81\% \pm 0.61\% \\ \mathbf{94.41\% \pm 0.10\%} \end{array}$	$\begin{array}{c} 77.07\% \pm 0.45\% \\ 77.50\% \pm 0.13\% \\ \textbf{77.52\% \pm 0.16\%} \end{array}$	1.00x 1.20x 1.19x
R50 R50 + 1D-SFD R50 + 1D-SWD	$\begin{array}{c} 93.31\% \pm 0.36\% \\ \textbf{93.85\%} \pm \textbf{0.40\%} \\ 93.67\% \pm 0.10\% \end{array}$	$\begin{array}{c} 76.18\% \pm 0.76\% \\ 77.38\% \pm 0.73\% \\ \textbf{77.46\% \pm 0.13\%} \end{array}$	1.00x 1.60x 1.25x
V16 V16 + 1D-SFD V16 + 1D-SWD	$\begin{array}{c} 93.27\% \pm 0.12\% \\ \textbf{93.49\%} \pm \textbf{0.07\%} \\ 93.42\% \pm 0.17\% \end{array}$	$\begin{array}{c} 72.48\% \pm 0.35\% \\ 72.27\% \pm 0.46\% \\ \textbf{72.57\%} \pm \textbf{0.32\%} \end{array}$	1.00x 1.17x 1.08x

improvements for all models compared to the baseline configurations. Specifically, the accuracy of ResNet50 on CIFAR-100 was improved by 1.28%pt. Additionally, integrating 2D-SWD on ResNet50 for CIFAR-100 improved performance by 1.8%pt.

Our analysis shows that on the CIFAR-10 dataset, both 1D and 2D Spectral Wavelet Dropout methods demonstrate performance that is competitive with, and sometimes even surpasses, that of 1D and 2D Spectral Fourier Dropout. For example, ResNet34 with 1D-SWD achieved an accuracy of  $94.41\% \pm 0.10\%$ , outperforming 1D and 2D-SFD.

While 2D Spectral Wavelet Dropout offers notable accuracy improvements, it also has a higher computational overhead compared to its 1D counterpart. Moreover, compared to the overheads associated with SFD methods, 1D-SWD demonstrates the lowest training overhead. This highlights the computational efficiency of 1D-SWD, making it a more practical choice for real-world applications where computational resources are constrained.

The regularization capability of 1D-SWD, demonstrated to be on par with or superior to 1D/2D-SFD. The combination of its computational efficiency and simplicity makes it an optimal choice for tasks where both performance and computational efficiency are valued.

## B. Image Classification on ImageNet

We extend our evaluation to the ImageNet dataset [17], which contains 1.2 million training images, 50,000 validation images, and 150,000 test images across 1,000 categories. We performed our experiments on ResNet50 across four distinct seeds, training each model for 450 epochs. Following the common practice, we report the top-1 classification accuracy on the validation set.

Considering the computational efficiency of 1D-SWD and the established presence of 2D-SFD in literature, our focus is on these methods. We have focused our scope due to the computational overhead associated with 1D-SFD and 2D-SWD. Specifically, 1D-SFD leads to a training time multiplier of 1.77x, while 2D-SWD extends this to 3.12x on ImageNet.

TABLE II 2D Spectral Dropout Methods on CIFAR-10/100: Accuracy and Training Time Multiplier (TTM).

Method	CIFAR-10 Acc.	CIFAR-100 Acc.	TTM
R18 R18 + 2D-SFD R18 + 2D-SWD	$\begin{array}{c} 94.04\% \pm 0.08\% \\ \mathbf{94.29\%} \pm \mathbf{0.19\%} \\ 94.22\% \pm 0.16\% \end{array}$	$\begin{array}{c} 76.47\% \pm 0.20\% \\ 77.21\% \pm 0.31\% \\ \textbf{77.22\% \pm 0.15\%} \end{array}$	1.00x 1.59x 2.22x
R34 R34 + 2D-SFD R34 + 2D-SWD	$\begin{array}{c} 93.69\% \pm 0.30\% \\ 93.97\% \pm 0.26\% \\ \mathbf{94.03\% \pm 0.31\%} \end{array}$	$\begin{array}{c} 77.07\% \pm 0.45\% \\ \textbf{77.49\% \pm 0.28\%} \\ 77.39\% \pm 0.36\% \end{array}$	1.00x 1.30x 2.21x
R50 R50 + 2D-SFD R50 + 2D-SWD	$\begin{array}{c} 93.31\% \pm 0.36\% \\ 93.59\% \pm 0.40\% \\ \textbf{93.62\% \pm 0.31\%} \end{array}$	$\begin{array}{c} 76.18\% \pm 0.76\% \\ 77.63\% \pm 0.55\% \\ \textbf{77.98\% \pm 0.10\%} \end{array}$	1.00x 1.87x 2.91x
V16 V16 + 2D-SFD V16 + 2D-SWD	$\begin{array}{c} 93.27\% \pm 0.12\% \\ 93.48\% \pm 0.12\% \\ \textbf{93.50\% \pm 0.21\%} \end{array}$	$\begin{array}{c} 72.48\% \pm 0.35\% \\ 72.52\% \pm 0.46\% \\ \textbf{72.63\% \pm 0.11\%} \end{array}$	1.00x 1.12x 1.88x

TABLE III Comparison of Dropout Techniques on ImageNet: Top-1 Accuracy and Training Time Multiplier (TTM).

Method	Top-1 Accuracy	TTM
ResNet50 (Baseline)	$76.87\% \pm 0.15\%$	1.00x
ResNet50 + D2D	$77.27\% \pm 0.12\%$	1.13x
ResNet50 + DB	$77.40\% \pm 0.11\%$	1.21x
ResNet50 + 1D-SWD	${f 77.47\% \pm 0.05\%}$	1.28x
ResNet50 + 2D-SFD	${f 77.87\%\pm 0.06\%}$	2.04x

We also include a comparison with modern dropout techniques such as dropblock (DB) and dropout2d (D2D) [10]. Comprehensive implementation details are provided in the Appendix, Section VII-E.

1) Results: Table III shows the effectiveness of our approach on ImageNet. For ResNet50, the application of 1D-SWD resulted in an accuracy increase of 0.60%pt compared to the baseline. It surpasses both DB and D2D. For ResNet50, the application of 1D-SWD resulted in an accuracy increase of 0.60%pt.

While 2D-SFD shows the highest performance gain, it requires tuning two hyperparameters, which is not needed with 1D-SWD. Additionally, the runtime overhead of 2D-SFD, noted as more than double (2.04 times) the baseline for ResNet50 on ImageNet, is significantly higher than the more manageable increase of approximately 1.28 times overhead of 1D-SWD. This difference not only highlights the computational efficiency of 1D-SWD but also its attractiveness as a simplified alternative to 2D-SFD.

## C. Object Detection on Pascal VOC

We compared SWD in both 1D and 2D forms to the established 2D-SFD and its one-dimensional counterpart on object detection. We tested these methods on the PASCAL VOC dataset [32], which includes a diverse set of 20 object categories. We use the mmobjectdetection library [35] to train our models.

We used the Faster R-CNN framework [2] and integrated our Spectral Wavelet Dropout approach to evaluate its impact

#### TABLE IV

PERFORMANCE OF SPECTRAL DROPOUT METHODS IN OBJECT DETECTION ON PASCAL VOC: MEAN AVERAGE PRECISION (MAP), AVERAGE PRECISION AT 50% IOU (AP50), AND TRAINING TIME MULTIPLIER (TTM). THE TRAINING OVERHEAD IS PRESENTED AS A FACTOR OF THE BASELINE MODEL'S COMPUTATIONAL TIME, WHERE "1.00X" REPRESENTS THE BASELINE.

Method	mAP (%)	AP50 (%)	TTM
FRCNN (Baseline)	$77.61\% \pm 0.27\%$	$77.62\% \pm 0.26\%$	1.00x
FRCNN + 1D-SFD FRCNN + 1D-SWD	$\begin{array}{c} 77.70\% \pm 0.17\% \\ \mathbf{78.01\% \pm 0.16\%} \end{array}$	$\begin{array}{c} 77.68\% \pm 0.17\% \\ \mathbf{78.00\%} \pm \mathbf{0.14\%} \end{array}$	3.39x 1.58x
FRCNN + 2D-SFD FRCNN + 2D-SWD	$\begin{array}{c} 77.86\% \pm 0.11\% \\ \textbf{78.17\%} \pm \textbf{0.18\%} \end{array}$	$\begin{array}{c} 77.86\% \pm 0.10\% \\ \mathbf{78.16\% \pm 0.19\%} \end{array}$	4.07x 3.51x

on model performance. We selected Mean Average Precision (mAP) and Average Precision at 50% Intersection over Union (AP50) as our main metrics. Further implementation details can be found in the Appendix, Section VII-F.

Table IV shows the experimental results. The 2D-SWD variant achieved the highest mAP and AP50 scores, demonstrating its superiority over SFD in object detection. Notably, 1D-SWD also showed a performance improvement while having much lower training overhead compared to 1D/2D-SFD and 2D-SWD. Both 1D-SWD and 2D-SWD outperformed their 1D and 2D Spectral Fourier Dropout counterparts. This shows the effectiveness of SWD as an alternative regularization method for object detection.

#### V. ANALYSIS AND DISCUSSION

#### A. Hyperparameters

For detailed outcomes of the hyperparameter optimization on CIFAR-10 and CIFAR-100, refer to Tables VIII and IX in Appendix VII-H. Results for ImageNet are presented in Table X. For SWD it is preferred to use a small dropout rate, i.e.  $p = \{0.1, 0.2\}$ . Increasing the dropout rate resulted in a decrease in the overall performance. In contrast to SWD, SFD needs different dropout and pruning rates for each setting.

## B. Frequency Dropout - 1D vs 2D

The decision to use one-dimensional (1D) or twodimensional (2D) Spectral Dropout, which includes both Wavelet and Fourier approaches, depends on their decomposition strategy. 1D methods target horizontal features, which are effectively for tasks with dominant horizontal patterns. Conversely, 2D techniques analyze both vertical and horizontal dimensions, therefore enhancing feature extraction at the cost of increased computational demand. Our findings show a trade-off between computational efficiency and performance enhancement that is influenced by the complexity of the task. For less demanding tasks, where the trade-off is minimal, we recommend the use of 1D Spectral Dropout due to its efficiency. In more challenging the choice between 2D Spectral Dropout for optimal performance and 1D Spectral Dropout as a more computationally efficient regularization method must be considered carefully.

## C. Computational Complexity Analysis

We compare the computational complexities and runtime performances of Spectral Wavelet Dropout and Spectral Fourier Dropout. Let n be the number of rows/columns of a square matrix. Theoretically, 2D-SFD and 1D-SFD have complexities of  $\mathcal{O}(n^2 \log n)$  for frequency transformations, indicating higher computational demands as the feature map sizes increase [36]. Conversely, SWD, in both 1D and 2D forms, operates at a reduced complexity of  $\mathcal{O}(n^2)$ . This improves efficiency due to a direct correlation with the pixel count.

Details on the runtime overheads can be found in the "Train Time Multiplier" column across Tables I, II, III, and IV. The runtime overhead of the 2D-SWD variant exceed that of 2D Spectral Fourier Dropout, likely due to less optimized wavelet transform algorithms or suboptimal GPU utilization. Notably, the Fourier Transform, known for its higher theoretical computational demands, showed improved efficiency on smaller datasets. This is attributed to the utilization of optimized GPUaccelerated libraries.

Our findings show that 1D-SWD consistently has a lower runtime overheads than both 1D and 2D-SFD on image classification and object detection task. Note that the runtime overhead by SWD is limited to the training phase and does not affect the inference time.

## D. Dropping Fourier Coefficients vs Frequency Bands

A main difference between SWD and SFD is the type of the dropped values. Specifically, SWD targets frequency bands of a wavelet transformed feature map, while SFD drops random Fourier coefficients in a Fourier-transformed feature map.

An advantage of dropping specific frequency bands is the possibility to control the level of regularization. By choosing specific frequency bands of the wavelet decomposition and adjusting the dropout rate, one can fine-tune the regularization strength based on the characteristics of the task and the desired model complexity.

In contrast to the Fourier transform, the Wavelet transform allows a multi-resolution analysis. Dropping the frequency bands forces the model to use information at different frequency scales.

Moreover, not all Fourier coefficients contribute equally to the information content of the image. Randomly dropping coefficients may unintentionally remove important frequency components, leading to a potential loss of crucial information. Furthermore, such random elimination could also have no effect at all, if the Fourier coefficients do not contribute to the information content of the image.

## VI. CONCLUSION AND FUTURE WORK

In this work, we introduced Spectral Wavelet Dropout, a novel regularization technique designed to improve the generalization capabilities of CNNs by randomly dropping frequency bands in wavelet-transformed feature maps. Spectral Wavelet Dropout is presented in two variants: 1D-SWD and 2D-SWD, which use different wavelet transformations. Empirical evaluations show that Spectral Wavelet Dropout matches and in some cases surpasses Spectral Fourier Dropout on datasets such as CIFAR-10, CIFAR-100 and ImageNet. Notably, the 1D variant of SWD achieves this competitive performance with a significant reduction in computational complexity, distinguishing it from both its 2D counterpart and SFD alternatives. In addition, SWD simplifies optimization by using just one hyperparameter, unlike the two required by SFD. For object detection tasks on the Pascal VOC dataset, both the 1D and 2D variants of Spectral Wavelet Dropout perform better than their Spectral Fourier Dropout counterparts while requiring less computation. Moreover, SWD requires only one hyperparameter instead of two in SFD.

While our current investigation has shown promising results, the full potential of SWD remains to be fully unlocked. Future research directions include the investigation of different wavelet functions, exploration of adaptive wavelets, and deeper analysis of decomposition levels to further refine the efficacy of SWD. Additionally, optimizing the handling of border distortions and exploring selective regularization based on frequency band energy distribution could provide avenues for enhancing the performance and applicability of SWD.

#### REFERENCES

- K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in 2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016. IEEE Computer Society, 2016, pp. 770–778. [Online]. Available: https://doi.org/10.1109/CVPR.2016.90
- [2] S. Ren, K. He, R. B. Girshick, and J. Sun, "Faster R-CNN: towards real-time object detection with region proposal networks," in Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada, C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, Eds., 2015, pp. 91– 99. [Online]. Available: https://proceedings.neurips.cc/paper/2015/hash/ 14bfa6bb14875e45bba028a21ed38046-Abstract.html
- [3] L. Wan, M. D. Zeiler, S. Zhang, Y. LeCun, and R. Fergus, "Regularization of Neural Networks using DropConnect," in *Proceedings of the 30th International Conference on Machine Learning*, *ICML 2013, Atlanta, GA, USA, 16-21 June 2013*, ser. JMLR Workshop and Conference Proceedings, vol. 28. JMLR.org, 2013, pp. 1058–1066. [Online]. Available: http://proceedings.mlr.press/v28/wan13.html
- [4] I. J. Goodfellow, Y. Bengio, and A. C. Courville, *Deep Learning*, ser. Adaptive computation and machine learning. MIT Press, 2016. [Online]. Available: http://www.deeplearningbook.org/
- [5] R. Tibshirani, "Regression Shrinkage and Selection via the Lasso," *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 58, no. 1, pp. 267–288, 1996. [Online]. Available: http: //www.jstor.org/stable/2346178
- [6] A. Krogh and J. A. Hertz, "A Simple Weight Decay Can Improve Generalization," in Advances in Neural Information Processing Systems 4, [NIPS Conference, Denver, Colorado, USA, December 2-5, 1991], J. E. Moody, S. J. Hanson, and R. Lippmann, Eds. Morgan Kaufmann, 1991, pp. 950–957. [Online]. Available: http://papers.nips.cc/paper/ 563-a-simple-weight-decay-can-improve-generalization
- [7] C. Shorten and T. M. Khoshgoftaar, "A survey on Image Data Augmentation for Deep Learning," J. Big Data, vol. 6, p. 60, 2019. [Online]. Available: https://doi.org/10.1186/s40537-019-0197-0
- [8] D. W. Opitz and R. Maclin, "Popular Ensemble Methods: An Empirical Study," J. Artif. Intell. Res., vol. 11, pp. 169–198, 1999. [Online]. Available: https://doi.org/10.1613/jair.614
- [9] N. Srivastava, G. E. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *J. Mach. Learn. Res.*, vol. 15, no. 1, pp. 1929–1958, 2014. [Online]. Available: http://dl.acm.org/citation.cfm?id=2670313

- [10] G. Ghiasi, T.-Y. Lin, and Q. V. Le, "DropBlock: A regularization method for convolutional networks," in Advances in Neural Information Processing Systems, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds., vol. 31. Curran Associates, Inc., 2018. [Online]. Available: https://proceedings.neurips.cc/paper/ 2018/file/7edcfb2d8f6a659ef4cd1e6c9b6d7079-Paper.pdf
- [11] J. Tompson, R. Goroshin, A. Jain, Y. LeCun, and C. Bregler, "Efficient object localization using Convolutional Networks," in *IEEE Conference* on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015. IEEE Computer Society, 2015, pp. 648–656. [Online]. Available: https://doi.org/10.1109/CVPR.2015.7298664
- [12] O. Rippel, J. Snoek, and R. P. Adams, "Spectral Representations for Convolutional Neural Networks," in Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada, C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, Eds., 2015, pp. 2449– 2457. [Online]. Available: https://proceedings.neurips.cc/paper/2015/ hash/536a76f94cf7535158f66cfbd4b113b6-Abstract.html
- [13] S. H. Khan, M. Hayat, and F. Porikli, "Regularization of Deep Neural Networks with Spectral Dropout," *Neural Networks*, vol. 110, pp. 82–90, 2019. [Online]. Available: https://doi.org/10.1016/j.neunet.2018.09.009
- [14] T. L. Williams and R. Li, "Advanced Image Classification Using Wavelets and Convolutional Neural Networks," in 15th IEEE International Conference on Machine Learning and Applications, ICMLA 2016, Anaheim, CA, USA, December 18-20, 2016. IEEE Computer Society, 2016, pp. 233–239. [Online]. Available: https: //doi.org/10.1109/ICMLA.2016.0046
- [15] J. Jo and Y. Bengio, "Measuring the tendency of cnns to learn surface statistical regularities," *CoRR*, vol. abs/1711.11561, 2017. [Online]. Available: http://arxiv.org/abs/1711.11561
- [16] R. Geirhos, P. Rubisch, C. Michaelis, M. Bethge, F. A. Wichmann, and W. Brendel, "ImageNet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness," in 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019. OpenReview.net, 2019. [Online]. Available: https://openreview.net/forum?id=Bygh9j09KX
- [17] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. S. Bernstein, A. C. Berg, and L. Fei-Fei, "ImageNet Large Scale Visual Recognition Challenge," *Int. J. Comput. Vis.*, vol. 115, no. 3, pp. 211–252, 2015. [Online]. Available: https://doi.org/10.1007/s11263-015-0816-y
- [18] J. Kukacka, V. Golkov, and D. Cremers, "Regularization for Deep Learning: A Taxonomy," *CoRR*, vol. abs/1710.10686, 2017. [Online]. Available: http://arxiv.org/abs/1710.10686
- [19] T. Devries and G. W. Taylor, "Improved Regularization of Convolutional Neural Networks with Cutout," *CoRR*, vol. abs/1708.04552, 2017. [Online]. Available: http://arxiv.org/abs/1708.04552
- [20] S. E. Finder, Y. Zohav, M. Ashkenazi, and E. Treister, "Wavelet Feature Maps Compression for Image-to-Image CNNs," *CoRR*, vol. abs/2205.12268, 2022. [Online]. Available: https://doi.org/10.48550/ arXiv.2205.12268
- [21] D. S. Taubman and M. W. Marcellin, JPEG2000 image compression fundamentals, standards and practice, ser. The Kluwer international series in engineering and computer science. Kluwer, 2002, vol. 642. [Online]. Available: https://doi.org/10.1007/978-1-4615-0799-4
- [22] S. Fujieda, K. Takayama, and T. Hachisuka, "Wavelet Convolutional Neural Networks," *CoRR*, vol. abs/1805.08620, 2018. [Online]. Available: http://arxiv.org/abs/1805.08620
- [23] P. Liu, H. Zhang, W. Lian, and W. Zuo, "Multi-Level Wavelet Convolutional Neural Networks," *IEEE Access*, vol. 7, pp. 74973– 74985, 2019. [Online]. Available: https://doi.org/10.1109/ACCESS. 2019.2921451
- [24] T. L. Williams and R. Li, "Wavelet Pooling for Convolutional Neural Networks," in 6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 -May 3, 2018, Conference Track Proceedings. OpenReview.net, 2018. [Online]. Available: https://openreview.net/forum?id=rkhlb8lCZ
- [25] M. Wolter and J. Garcke, "Adaptive wavelet pooling for convolutional neural networks," in *The 24th International Conference on Artificial Intelligence and Statistics, AISTATS 2021, April 13-15, 2021, Virtual Event*, ser. Proceedings of Machine Learning Research, A. Banerjee and K. Fukumizu, Eds., vol. 130. PMLR, 2021, pp. 1936–1944. [Online]. Available: http://proceedings.mlr.press/v130/wolter21a.html

- [26] M. Wolter, S. Lin, and A. Yao, "Neural Network Compression via Learnable Wavelet Transforms," in Artificial Neural Networks and Machine Learning - ICANN 2020 - 29th International Conference on Artificial Neural Networks, Bratislava, Slovakia, September 15-18, 2020, Proceedings, Part II, ser. Lecture Notes in Computer Science, I. Farkas, P. Masulli, and S. Wermter, Eds., vol. 12397. Springer, 2020, pp. 39–51. [Online]. Available: https://doi.org/10.1007/978-3-030-61616-8\_4
- [27] A. Haar, "Zur theorie der orthogonalen funktionensysteme. (erste mitteilung)." *Mathematische Annalen*, vol. 69, pp. 331–371, 1910. [Online]. Available: http://eudml.org/doc/158469
- [28] I. Daubechies, *Ten Lectures on Wavelets*. Society for Industrial and Applied Mathematics, 1992. [Online]. Available: https://epubs.siam.org/ doi/abs/10.1137/1.9781611970104
- [29] M. Mohlenkamp, "A Tutorial on Wavelets and Their Applications," 2004. [Online]. Available: http://www.ohiouniversityfaculty.com/ mohlenka/20044/PASIII/waveletIPAM.pdf
- [30] C. Shannon, "Communication in the Presence of Noise," *Proceedings of the IRE*, vol. 37, no. 1, pp. 10–21, 1949.
- [31] A. Krizhevsky and G. Hinton, "Learning multiple layers of features from tiny images," *Master's thesis, Department of Computer Science, University of Toronto*, 2009.
- [32] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The pascal visual object classes (voc) challenge," *International Journal of Computer Vision*, vol. 88, no. 2, pp. 303–338, Jun. 2010.
- [33] F. Cotter, "Uses of complex wavelets in deep convolutional neural networks," Ph.D. dissertation, University of Cambridge, UK, 2019. [Online]. Available: https://ethos.bl.uk/OrderDetails.do?uin=uk.bl.ethos. 810067
- [34] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," in 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings, Y. Bengio and Y. LeCun, Eds., 2015. [Online]. Available: http://arxiv.org/abs/1409.1556
- [35] K. Chen, J. Wang, J. Pang, Y. Cao, Y. Xiong, X. Li, S. Sun, W. Feng, Z. Liu, J. Xu, Z. Zhang, D. Cheng, C. Zhu, T. Cheng, Q. Zhao, B. Li, X. Lu, R. Zhu, Y. Wu, J. Dai, J. Wang, J. Shi, W. Ouyang, C. C. Loy, and D. Lin, "MMDetection: Open mmlab detection toolbox and benchmark," arXiv preprint arXiv:1906.07155, 2019.
- [36] H. Guo and C. Burrus, "Wavelet transform based fast approximate Fourier transform," in 1997 IEEE International Conference on Acoustics, Speech, and Signal Processing, vol. 3, 1997, pp. 1973–1976 vol.3.
- [37] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Z. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "PyTorch: An Imperative Style, High-Performance Deep Learning Library," in Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada, H. M. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. B. Fox, and R. Garnett, Eds., 2019, pp. 8024–8035. [Online]. Available: https://proceedings.neurips. cc/paper/2019/hash/bdbca288fee7f92f2bfa9f7012727740-Abstract.html
- [38] K. He, X. Zhang, S. Ren, and J. Sun, "Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification," in 2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015. IEEE Computer Society, 2015, pp. 1026–1034. [Online]. Available: https://doi.org/10.1109/ICCV.2015. 123

## VII. APPENDIX

## A. Visualizing 1D Spectral Wavelet Dropout

Figure 3 illustrates the 1D Spectral Wavelet Dropout applied to a feature map  $X \in \mathbb{R}^{C \times H \times W}$ . The process begins by flattening the feature map along each channel, resulting in  $\hat{X} \in \mathbb{R}^{C \times H \cdot W}$ . Then a three-level one-dimensional discrete wavelet decomposition (1D-DWT) is performed on  $\hat{X}$ , resulting in approximation and detailed frequency coefficients across three levels. They are also referred as frequency bands. In the next phase, 1D-SWD randomly drops detailed wavelet bands using a dropout probability p. This introduces regularization. This regularization is shown in Figure 3 by the dropping of the Level 2 coefficients. The feature map is transformed back into the spatial domain using an inverse 1D-DWT. Finally, it is reshaped for input into the subsequent network layer.

#### B. Implementing the 2D Spectral Wavelet Dropout Algorithm

At the core of the 2D-SWD, as detailed in Algorithm 2, is the application of the 2D Discrete Wavelet Transform (2D-DWT). This transformation process decomposes each channel of the input feature map into four distinct sub-bands: LL (low-low), LH (low-high), HL (high-low), and HH (high-high). These represent the approximation, horizontal detail, vertical detail, and diagonal detail components.

Following the decomposition, a dropout mask is sampled using a dropout probability p. This mask is applied to the detail coefficients (LH, HL), and HH. This drops a portion of the wavelet coefficients to introduce regularization. The approximation coefficients (LL) are preserved to maintain the overall structure of the image. Scaling the rest of the coefficients by a factor of  $(1 - p)^{-1}$  compensates for the dropout's reduction of signal energy.

The final step involves transforming the feature map back into the spatial domain using the inverse 2D-DWT.

_	
	Algorithm 2: Spectral Wavelet Dropout - 2D-SWD
	<b>Parameter:</b> dropout probability = $p \in (0, 1)$
	<b>Input:</b> $X \in \mathbb{R}^{B \times C \times H \times W}$
	<b>Output:</b> $\hat{X} \in \mathbb{R}^{B \times C \times H \times W}$
	// Save size of $X$
1	size = $X$ .size()
	// Compute the 1-level 2D-DWT for
	each channel using the db3-wavelet
2	LL, LH, HL, HH = 2D-DWT(X, J = 1, wave='db3')
	// Create mask for dropout
3	dropout_mask = Bernoulli $(3, \text{prob} = 1 - p)$
	// Wavelet Dropout
4	$LH = LH \cdot dropout\_mask[0]/(1-p)$
5	$HL = HL \cdot dropout\_mask[1]/(1-p)$
6	$HH = HH \cdot dropout\_mask[2]/(1-p)$
	// Compute inverse Wavelet
	transformation
7	$\hat{X} = 2\text{D-IDWT}([LL, LH, HL, HH], \text{wave='db3'})$

## C. In-Depth Analysis of the Discrete Wavelet Transform in SWD

Figure 4 illustrates the frequency domain representation of the Discrete Wavelet Transform. It demonstrates the DWT's ability to partition a signal into multiple frequency components.

## D. Detailed CIFAR-10/100 Training Procedures

Our CIFAR experiments used PyTorch 1.10.1 [37] on an Nvidia GeForce 1080Ti GPU. We used the same training setup for all models, with a batch size of 128 and the SGD optimizer with a momentum of 0.9. The initial learning rate was set to 0.1, with adjustments based on the dataset: for CIFAR-10, it was reduced by a factor of 0.1 at epochs 90 and 136, along with a weight decay of 1e - 4 [1]; for CIFAR-100, it was reduced by 0.2 at epochs 60, 120, and 160, using a weight decay of 5e - 4. These parameter settings match those recommended in [1], without further hyperparameter optimization. Data augmentation techniques included random cropping (size 32 with padding 4), random horizontal flipping, and normalization, as outlined in [7], with normalization being the only preprocessing step for the test set. Network initialization followed the Kaiming uniform approach [38], with VGG's linear layers similarly initialized and its convolutional layers set to a Gaussian distribution ( $\mu = 0, \sigma = \sqrt{2/n}$ ).

#### E. Detailed ImageNet Training Procedures

In our experiments, we used PyTorch 1.10.1 [37] with four Nvidia GeForce 1080Ti GPUs. All models were trained with the SGD optimizer, set with a momentum of 0.9 and an initial learning rate of 0.1, which was reduced by a factor of 0.1 at epochs 75, 150, 225, 300, and 375. A weight decay factor of 1e - 4 was applied, and the batch size was standardized at 128.

For data augmentation, the training dataset was randomly cropped to  $224 \times 224$ , randomly flipped horizontally, and normalized. The validation dataset was resized to  $256 \times 256$ , center cropped to  $224 \times 224$ , and normalized. Our setup matches the configuration specified in the PyTorch ImageNet training example<sup>2</sup>.

## F. Implementation Details on PASCAL VOC using Faster R-CNN

In our object detection experiments on the PASCAL VOC dataset, we use the mmobjectdetection library [35] for setting up and training our models. This section outlines the key implementation details to ensure our results can be reproduced.

*a) Model Configuration:* We set up our Faster R-CNN [2] with a ResNet-50 backbone and a Feature Pyramid Network (FPN) for multi-scale feature extraction.

<sup>&</sup>lt;sup>2</sup>https://github.com/pytorch/examples/blob/main/imagenet

#### 1D-SWD



Fig. 3. Schematic of the 1D-SWD applied to a feature map  $X \in \mathbb{R}^{C \times H \times W}$ . The process begins by flattening X across channels, resulting in  $\hat{X} \in \mathbb{R}^{C \times H \cdot W}$ . A three-level 1D-DWT is then applied to  $\hat{X}$ , yielding approximation and detailed coefficients across three levels. For regularization, detailed coefficients are dropped using a dropout probability p, with Level 2 Coefficients being dropped in the example. The regularized feature map is transformed back using an inverse 1D-DWT.



Fig. 4. This figure shows the Discrete Wavelet Transform dividing a signal into distinct frequency bands. Such multi-resolution analysis is important for our Spectral Wavelet Dropout method, as it allows randomly dropping of information across different scales.

b) Data Preprocessing and Augmentation: Our preprocessing pipeline includes several steps to prepare images for object detection. First, images are loaded and annotations are retrieved. Then, images are resized to (1000, 600), keeping their original aspect ratio. To augment the data and introduce variability, we randomly flip images horizontally with a 50% probability. This augmentation is applied uniformly across the training dataset to enhance model robustness and generalization. For validation, images are resized in the same way but without random flips to keep evaluation consistent.

c) Training Configuration: The model is trained on the combined trainval sets of VOC2007 and VOC2012 and evaluated on the VOC2007 validation set. We use a batch size of 2. The SGD optimizer is used with momentum and weight decay, and the learning rate is adjusted according to a predefined schedule. We use the mean Average Precision (mAP) metric, calculated with the '11points' interpolation method, as the evaluation metric.

*d)* Spectral Dropout Integration: All Spectral Dropout methods were applied within the third block of the ResNet50 backbone.

TABLE V Performance on CIFAR-10 with 1D-SWD Targeting Level 3 Detail Coefficients.

Method	Accuracy
ResNet18 + 1D-SWD3	$94.05\% \pm 0.21\%$
ResNet34 + 1D-SWD3	$94.26\% \pm 0.23\%$
ResNet50 + 1D-SWD34	$94.00\% \pm 0.18\%$

TABLE VI Performance on CIFAR-100 with 1D-SWD Targeting Level 3 Detail Coefficients.

Method	Accuracy
ResNet18 + 1D-SWD4	$77.11\% \pm 0.21\%$
ResNet34 + 1D-SWD4	$77.32\%\pm 0.41\%$
ResNet50 + 1D-SWD4	$77.42\% \pm 0.47\%$

### G. Frequency Band Impact Analysis with 1D-SWD

We analyzed the effects of explicitly dropping different frequency components: the approximation coefficients, and level 1, 2, and 3 detail coefficients, both individually and in combinations. Our results are shown in Tables V, VI, and VII. The naming convention of model configurations, such as "ResNet50 + 1D-SWD34", indicates where 1D-SWD is applied within the ResNet50 architecture, showing the specific stages where frequency band manipulation happens.

They indicate a significant performance improvement when only level 3 detail coefficients are dropped across CIFAR-10/100 and ImageNet datasets. Our analysis shows that dropping approximation coefficients, which capture the lowfrequency, smoothed parts of the signal, is not effective. These coefficients contain important information about the signal's overall structure, whereas detail coefficients correspond to high-frequency variations and potential noise.

TABLE VII Performance on ImageNet with 1D-SWD Targeting Level 3 Detail Coefficients.

Method	Top-1 Accuracy
ResNet50 + 1D-SWD4	$77.51\%\pm0.11\%$

## H. Hyperparameters

1) Method Parameters: The hyperparameters for our CI-FAR and ImageNet experiments are shown in Tables VIII, IX, and X. We found that lower dropout rates, specifically  $p = \{0.1, 0.2\}$ , work best for both 1D and 2D Spectral Wavelet Dropout. Higher dropout rates led to lower performance.

In contrast, Spectral Fourier Dropout in both 1D and 2D forms required different dropout and pruning rates, showing the need for specific hyperparameter tuning for each setting.

Figures 5 and 6 show how ResNet50 performance on CIFAR-10 and CIFAR-100 varies with different hyperparameters for 2D-SFD. This underscores the importance of a comprehensive hyperparameter search for SFD.



Fig. 5. Hyperparameter search results for ResNet50 with 2D-SFD on CIFAR-10, illustrating the impact of various configurations on model performance.



Fig. 6. Hyperparameter search results for ResNet50 with 2D-SFD on CIFAR-100, demonstrating the performance variability across different hyperparameter settings.

2) Positional Analysis: We did an ablation study to investigate the effect of the position of 1D-SWD within the ResNets. We did the same ablation study for 1D-SFD, 2D-SWD and 2D-SFD. Similar to other dropping methods like DropBlock [10] it is preferred to insert the frequency dropout methods in deeper layers, e.g. for ResNets in the modules 3 or 4. The Tables XI

TABLE VIII Hyperparameters for CIFAR-10 and CIFAR-100 experiments with 1D spectral methods.

Model	<b>CIFAR-</b> 10	<b>CIFAR-</b> 100
ResNet18 + 1D-SWD	p = 0.1	p = 0.1
ResNet18 + 1D-SFD	$p = 0.1, \eta = 0.1$	$p = 0.2, \eta = 0.2$
ResNet34 + 1D-SWD	p = 0.2	p = 0.1
ResNet34 + 1D-SFD	$p = 0.3, \eta = 0.5$	$p = 0.1, \eta = 0.0$
ResNet50 + 1D-SWD	p = 0.2	p = 0.1
ResNet50 + 1D-SFD	$p = 0.2, \eta = 0.1$	$p = 0.3, \eta = 0.1$
VGG16 + 1D-SWD	p = 0.1	p = 0.2
VGG16 + 1D-SFD	$p = 0.3, \eta = 0.4$	$p = 0.4, \eta = 0.1$

TABLE IX Hyperparameters for CIFAR-10 and CIFAR-100 experiments with 2D spectral methods.

Model	<b>CIFAR-</b> 10	<b>CIFAR-</b> 100
ResNet18 + 2D-SWD ResNet18 + 2D-SFD	p = 0.1 $p = 0.1, \eta = 0.2$	$\begin{array}{c} p=0.2\\ p=0.1, \eta=0.1 \end{array}$
ResNet34 + 2D-SWD	p = 0.2	p = 0.2
ResNet34 + 2D-SFD	$p = 0.4, \eta = 0.3$	$p = 0.2, \eta = 0.0$
ResNet50 + 2D-SWD	p = 0.2	p = 0.1
ResNet50 + 2D-SFD	$p = 0.4, \eta = 0.4$	$p = 0.2, \eta = 0.0$
VGG16 + 2D-SWD	p = 0.5	p = 0.2
VGG16 + 2D-SFD	$p = 0.1, \eta = 0.4$	$p = 0.4, \eta = 0.5$

and XII show in which module the frequency dropout methods are inserted in our experiments.

The regularization methods can be too harsh when applied to higher resolution feature maps in earlier modules. Inserting the frequency dropout methods in the fourth module generally led to the best results. For smaller networks on easier tasks (e.g. ResNet18 on CIFAR10) it is preferred to insert them in the third module. This may be attributed due to the shallower architecture, which might need more regularization in the third module to prevent overfitting to noise or irrelevant details. In deeper ResNet variants, the increased model capacity enables them to capture finer details and patterns in the fourth module, making 1D-SWD more effective at this stage.

We also analyzed the effect of the position of 1D-SWD within the basic or bottleneck buildings blocks. We investigated inserting SWD before and after the convolution layers within the buildings blocks. Both ways increase the performance. However, the optimal position depends on the resolution of the feature maps.

For high-resolution images like ImageNet, 1D-SWD should be inserted before each convolution layer, following the sequence: 1D-SWD - Conv - BN - ReLU, including in the residual connection (1D-SWD - Conv - BN). This architecture forces the network to focus on various frequency bands.

For low-resolution images like CIFAR, it is preferred to place 1D-SWD after the convolution layers (Conv - 1D-SWD - BN - ReLU), both in the main and residual block. This allows the network to first extract features through convolution, which

 TABLE X

 Hyperparameters of experiments on ImageNet.

Model	ImageNet
ResNet50 + 1D-SWD	p = 0.2
ResNet50 + 2D-SFD	$p = 0.2, \eta = 0.1$

are then manipulated using 1D-SWD.

For both low- and high-resolution images, the 2D-SWD is strategically positioned before each convolutional layer to optimize performance.

TABLE XIPOSITION OF SPECTRAL DROPOUT ON CIFAR-10/100 - 1D.

Model	CIFAR-10	CIFAR-100
R18 + 1D-SFD	4	4
R18 + 1D-SWD	3	4
R34 + 1D-SFD	4	4
R34 + 1D-SWD	3	4
R50 + 1D-SFD	4	4
R50 + 1D-SWD	4	4
V16 + 1D-SWD	4	4

TABLE XIIPOSITION OF SPECTRAL DROPOUT ON CIFAR-10/100 - 2D.

Model	CIFAR-10	CIFAR-100
R18 + 2D-SFD	4	4
R18 + 2D-SWD	3	4
R34 + 2D-SFD	4	4
R34 + 2D-SWD	4	4
R50 + 2D-SFD	4	4
R50 + 2D-SWD	4	4
V16 + 2D-SWD	4	4