

On the Shift Invariance of Max Pooling Feature Maps in Convolutional Neural Networks

Anonymous authors
Paper under double-blind review

Abstract

This paper focuses on improving the mathematical interpretability of convolutional neural networks (CNNs) in the context of image classification. Specifically, we tackle the instability issue arising in their first layer, which tends to learn parameters that closely resemble oriented band-pass filters when trained on datasets like ImageNet. Subsampled convolutions with such Gabor-like filters are prone to aliasing, causing sensitivity to small input shifts. In this context, we establish conditions under which the max pooling operator approximates a complex modulus, which is nearly shift invariant. We then derive a measure of shift invariance for subsampled convolutions followed by max pooling. In particular, we highlight the crucial role played by the filter’s frequency and orientation in achieving stability. We experimentally validate our theory by considering a deterministic feature extractor based on the dual-tree complex wavelet packet transform, a particular case of discrete Gabor-like decomposition.

1 Introduction

Understanding the mathematical properties of deep convolutional neural networks (CNNs) (LeCun et al., 2015) remains a challenging issue today. On the other hand, wavelet and multi-resolution analysis are built upon a well-established mathematical framework. They have proven to be efficient for tasks such as signal compression and denoising (Vetterli, 2001), and have been widely used as feature extractors for signal, image and texture classification (Laine & Fan, 1993; Pittner & Kamarthi, 1999; Yen, 2000; Huang & Aviyente, 2008). There is a broad literature revealing strong connections between these two paradigms, as discussed in Sections 1.1 and 1.2. Inspired by this line of research, the present paper extends existing knowledge about CNN properties. Specifically, we assess the shift invariance of max pooling feature maps through both theoretical and empirical approaches in the context of image classification, by leveraging the properties of oriented band-pass filters.

1.1 Motivations and Main Contributions

CNNs process input images through convolutions and nonlinear pooling operations, transforming them into high-level feature vectors that are subsequently used for the task at hand. In image classification, the feature vectors are fed into a linear classifier. To achieve high classification accuracy, a convolutional network must preserve discriminative image features while reducing intra-class variability (LeCun et al., 1998; Bruna & Mallat, 2013). An important and often desired property of CNNs is their ability to remain invariant to small input transformations, such as translations, rotations, distortions, or scaling (Liao & Peng, 2010; Bruna & Mallat, 2013; Sifre & Mallat, 2013; Bietti & Mairal, 2017; Wiatowski & Bölcskei, 2018; Cahill et al., 2024).

In particular, an image in which the main subject is slightly shifted from its original position should retain its initial classification. The absence of this property, known as *translation invariance* or *shift invariance*, can negatively impact the model’s predictive performance. Consequently, a reliable model must learn this property through extensive training on large datasets. The main focus of this paper is therefore to assess whether shift invariance is inherently guaranteed by the model’s architecture. Since perfect invariance is rarely achieved, we also use the term *stability* to refer to this behavior.

More specifically, we focus on commonly observed phenomenon in CNNs when trained on image datasets: many convolution kernels in the first layer resemble band-pass oriented waveforms (Yosinski et al., 2014; Rai & Rivas, 2020), referred to as *Gabor-like filters*. Whether the features extracted by these filters remain stable under translations has been partly addressed by Azulay & Weiss (2019); Zhang (2019), who highlight that strided convolution and pooling operations can significantly diverge from shift invariance, due to aliasing effects when subsampling high-frequency signals. In response, Zhang (2019); Zou et al. (2023) introduced antialiasing methods based on low-pass filtering, improving both stability and predictive performance—albeit at the cost of some loss of information.

In the current paper, we show that, under specific conditions that we establish, the max pooling operator can partially restore shift invariance. We unveil a connection between the output of the first max pooling layer and the modulus of complex Gabor-like coefficients, which is known to be nearly shift invariant. This work offers a promising direction for improving shift invariance in CNNs while preserving high-frequency information—unlike the previously-mentioned approaches.

Before proceeding further, we emphasize that our study is not limited to purely convolutional architectures. In recent years, self-attention mechanisms have gained significant interest in computer vision due to their ability to model complex, long-range dependencies in image representations. Notably, the *vision transformer* (ViT) was introduced by Dosovitskiy et al. (2021) by adapting the transformer architecture, initially developed for natural language processing (NLP) (Vaswani et al., 2017), to computer vision tasks. Unlike CNNs, the ViT operates without any convolutional layers. Instead, input images are partitioned into fixed-size patches, which serve as inputs to the first self-attention module. However, more recent research has explored hybrid architectures that integrate self-attention with convolutional components (Wu et al., 2021; Yuan et al., 2021; Hassani et al., 2022; Li et al., 2023; Yin et al., 2024). This approach allows for reducing the amount of labeled data required, while achieving faster training and improving generalizability. In particular, the first layers of a CNN can be used as a “convolutional token embedding,” replacing the naive patch extraction used in the original ViT. The theoretical framework presented in this paper also applies to such hybrid architectures, providing a better understanding of the invariance properties of the inputs to self-attention modules.

1.2 Related Work

Analyzing the invariance properties of CNNs is critical as it enables to identify their shortcomings and provides an opportunity to enhance their performance. In recent years, several works focused on this topic.

1.2.1 Wavelet Scattering Networks

Most notably, Bruna & Mallat (2013) developed a family CNN-like architectures, named *wavelet scattering networks* (ScatterNets), based on a succession of complex convolutions with wavelet filters followed by nonlinear modulus pooling. They produce translation-invariant image representations which are stable to deformation and preserve high-frequency information (Mallat, 2012; 2016; Czaja et al., 2024). A variation has been proposed by Sifre & Mallat (2013) to include rotational invariance. ScatterNets achieve strong performance on handwritten digits and texture datasets, but do not scale well to more complex ones. To overcome this, Oyallon (2017); Oyallon et al. (2018) introduced hybrid ScatterNets, where the scattering coefficients are fed into a standard CNN architecture, showing that the network complexity can be reduced while keeping competitive performance. Derived models include ScatterNets built upon the dual-tree complex wavelet transform (Singh & Kingsbury, 2017), learnable and parametric ScatterNets (Cotter & Kingsbury, 2019; Gauthier et al., 2022), geometric ScatterNets operating on Riemannian manifolds (Perlmutter et al., 2020), and graph ScatterNets (Gama et al., 2019; Zou & Lerman, 2020). Also worth mentioning, Czaja & Li (2019; 2020) studied ScatterNets based on uniform covering frames, i.e., frames splitting the frequency domain into windows of roughly equal size, much like DT-CWPT frames (as used in the present paper). Other works by Zarka et al. (2020; 2021) proposed to sparsify wavelet scattering coefficients by learning a dictionary matrix, to learn 1×1 convolutions between feature maps of scattering coefficients and to apply soft thresholding to reduce within-class variability.

ScatterNets are specifically designed to meet some desired properties. As deep learning architectures with well-established mathematical properties, they are sometimes used as explanatory models for standard, freely-trained networks. However, whether their properties are transferable to a broader class of models is unclear, because the former rely on complex-valued convolutions whereas more conventional architectures exclusively employ real-valued kernels. Moreover, the modulus operator is used as an activation and pooling layer in ScatterNets, whereas standard CNNs implement pointwise nonlinear operators such as ReLU and spatial pooling layers such as max pooling. This limitation has been pointed out by Tygert et al. (2016) as an argument in favor of complex-valued CNNs. In this context, our work seeks evidence that properties established for complex-valued networks are—to some extent—embedded in standard architectures.

1.2.2 Invariance Studies in CNNs

Wiatowski & Bölcskei (2018) considered a wide variety of feature extractors involving convolutions, Lipschitz-continuous non-linearities and pooling operators. The paper shows that outputs become more translation invariant with increasing network depth. Additionally, Cahill et al. (2024) designed a family of operators called *max filters*, which encompass a wide variety of operators including, in specific cases, the max pooling operator. Stability with respect to diffeomorphisms were established, following the ideas developed for scattering networks. However, these results do not fully extend to the discrete framework, because subsampled convolutions with band-pass real-valued filters can introduce aliasing artifacts, resulting in instability to translations (Azulay & Weiss, 2019; Zhang, 2019). The current paper specifically addresses this issue.

Another line of work is focused on modeling and studying CNNs from the point of view of convolutional kernel networks (Bietti & Mairal, 2019a;b; Scetbon & Harchaoui, 2020; Bietti, 2022; Chen, 2023). These authors showed that certain classes of CNNs are contained in the reproducing kernel Hilbert space (RKHS) of a multilayer convolutional kernel representation. As such, stability metrics are estimated, based on the RKHS norm which is difficult to control in practice. Kernel representations do not seem to suffer from aliasing effects; this can be explained by the Gaussian pooling layers that have been employed instead of max pooling: by discarding high-frequency information, shift invariance is preserved.

Finally, some papers studied stability of CNNs in a broader sense, measured in terms of Lipschitz continuity (Szegedy et al., 2014; Balan et al., 2018; Virmaux & Scaman, 2018; Pérez et al., 2020; Zou et al., 2020; Gupta et al., 2022; Zühlke & Kudenko, 2024). However, the Lipschitz bounds, which have been obtained theoretically, are generally several orders of magnitude higher than empirical results. This discrepancy may be due to the fact that these bounds were obtained for generic situations and represent overly conservative worst-case scenarios, rather than typical real-world situations. Furthermore, the specific case of convolutions with band-pass Gabor-like filters have been overlooked, except for Pérez et al. (2020).

In summary, we have identified the following blind spots in the literature, regarding the topic of studying shift invariance in CNNs.

- The effect of the max pooling operator on network stability under small input shifts has not been investigated, particularly when used in combination with Gabor-like convolutions.
- While the shift invariance of CNNs tends to increase with network depth in the continuous framework, in the discrete case, the presence of subsampled convolutions with oriented band-pass filters can lead to aliasing artifacts. To our knowledge, the literature lacks theoretical studies that take these aliasing effects into account.
- Although extensive studies have been conducted on complex-valued convolutions followed by modulus, a link is missing to extend these results to standard CNNs, which implement real-valued convolutions and spatial pooling operators.

All these points have been tackled in the present paper, from both theoretical and empirical perspectives.

1.3 Paper Outline

In what follows, $l_{\mathbb{R}}^2(\mathbb{Z}^2)$ and $l_{\mathbb{C}}^2(\mathbb{Z}^2)$ represent the discrete spaces of square-summable two-dimensional sequences with values in \mathbb{R} and \mathbb{C} , respectively. Let $W \in l_{\mathbb{C}}^2(\mathbb{Z}^2)$ denote a two-dimensional band-pass, oriented and analytic *Gabor-like* filter, for which a formal definition will be provided in (11). We first consider an operator, referred to as *real-max-pooling* ($\mathbb{R}\text{Max}$), which computes the subsampled cross-correlation between an input image $X \in l_{\mathbb{R}}^2(\mathbb{Z}^2)$ and the real part of W ; then calculates the maximum value over a sliding discrete grid:

$$U_{m,q}^{\max}[W] : X \mapsto \text{MaxPool}_q((X * \overline{\text{Re } W}) \downarrow m), \quad (1)$$

where $m \in \mathbb{N} \setminus \{0\}$ denotes a subsampling factor, \bar{V} denotes the “flipped” sequence for any given $V \in l_{\mathbb{R}}^2(\mathbb{Z}^2)$ or $l_{\mathbb{C}}^2(\mathbb{Z}^2)$, satisfying, for any $\mathbf{n} \in \mathbb{Z}^2$,

$$\bar{V}[\mathbf{n}] := V[-\mathbf{n}], \quad (2)$$

and $*$, \downarrow respectively refer to the convolution and subsampling operations, defined by

$$(X * \bar{V})[\mathbf{n}] := \sum_{\mathbf{p} \in \mathbb{Z}^2} X[\mathbf{p}] \bar{V}[\mathbf{n} - \mathbf{p}] \quad \text{and} \quad (Y \downarrow m)[\mathbf{n}] := Y[m\mathbf{n}]. \quad (3)$$

In the above expression, MaxPool_q selects the maximum value over a sliding grid of size $(2q+1) \times (2q+1)$, with a subsampling factor of 2. More formally, for any $Y \in l_{\mathbb{R}}^2(\mathbb{Z}^2)$ and any $\mathbf{n} \in \mathbb{Z}^2$,

$$\text{MaxPool}_q(Y)[\mathbf{n}] := \max_{\|\mathbf{p}\|_{\infty} \leq q} Y[2\mathbf{n} + \mathbf{p}]. \quad (4)$$

On the other hand, we consider an operator, referred to as *complex-modulus* ($\mathbb{C}\text{Mod}$), computing the modulus of subsampled cross-correlation between X and W :

$$U_m^{\text{mod}}[W] : X \mapsto |(X * \bar{W}) \downarrow (2m)|. \quad (5)$$

First, we show that, under the Gabor hypothesis, $\mathbb{C}\text{Mod}$ is stable with respect to small input shifts (Section 2). We then establish conditions on the filter’s frequency and orientation under which $\mathbb{C}\text{Mod}$ and $\mathbb{R}\text{Max}$ produce comparable outputs (Section 3):

$$U_m^{\text{mod}}[W](X) \approx U_{m,q}^{\max}[W](X). \quad (6)$$

We deduce a measure of shift invariance for $\mathbb{R}\text{Max}$ operators, which benefits from the stability of $\mathbb{C}\text{Mod}$ (Section 4). Next, we extend our results to multichannel operators (*i.e.*, applied on RGB input images), such as implemented in conventional CNN architectures (Section 5). Our framework therefore provides a theoretical grounding to study these networks.

Remark 1. In the above definitions, cross-correlations are computed with a subsampling factor which is twice larger for $\mathbb{C}\text{Mod}$, compared to $\mathbb{R}\text{Max}$. However, since max pooling is also computed with subsampling, both operators have the same subsampling factor of $2m$.

Finally, in Section 6, we assess our theoretical findings on a deterministic setting based on the dual-tree complex wavelet packet transform (DT-CWPT), a particular case of discrete Gabor-like decomposition with perfect reconstruction properties (Bayram & Selesnick, 2008). DT-CWPT spawns a set of convolution kernels which tile the Fourier domain into square regions of identical size. Such kernels possess characteristics that are comparable to those found in the first convolution layer of CNNs after training with image datasets such as ImageNet (Russakovsky et al., 2015). More specifically, given an input image, we compute the mean square error between the outputs of $\mathbb{C}\text{Mod}$ and $\mathbb{R}\text{Max}$, for each wavelet packet filter. We then observe that shift invariance, when measured on $\mathbb{R}\text{Max}$ feature maps, is nearly achieved when they remain close to $\mathbb{C}\text{Mod}$ outputs. We therefore establish a domain of validity for shift invariance of the $\mathbb{R}\text{Max}$ operator.

This work builds upon an idea sketched by Waldspurger (2015, pp. 190–191), which suggests a potential connection between the combinations “real wavelet transform \rightarrow max pooling” on the one hand and “complex wavelet transform \rightarrow modulus” on the other hand. Building on this idea, we investigated whether the invariance properties of complex moduli could be captured by the max pooling operator. However, as shown in this paper, these principles do not fully extend to the discrete framework. To address this limitation, we adopted a Bayesian point of view.

2 Shift Invariance of CMod Outputs

The primary goal of this paper is to theoretically establish conditions for near-shift invariance at the output of the first max pooling layer. In this section, we start by proving shift invariance of CMod operators. Then, in Section 3, we establish conditions under which RMax and CMod produce closely related outputs. Finally, in Section 4, we derive a probabilistic measure of shift invariance for RMax.

2.1 Notations

The complex conjugate of any number $z \in \mathbb{C}$ is denoted by z^* . For any $p \in \mathbb{R}_{>0} \cup \{\infty\}$, $\mathbf{x} \in \mathbb{R}^2$ and $r \in \mathbb{R}_+$, we denote by $B_p(\mathbf{x}, r) \subset \mathbb{R}^2$ the closed l^p -ball with center \mathbf{x} and radius r . When $\mathbf{x} = \mathbf{0}$, we write $B_p(r)$.

Continuous Framework Considering a measurable subset E of \mathbb{R}^2 , we denote by $L_{\mathbb{C}}^2(E)$ the Hilbert space of square-integrable functions $F : E \rightarrow \mathbb{C}$. Whenever we talk about equality in $L_{\mathbb{C}}^2(E)$ or inclusion in E , it shall be understood as “almost everywhere with respect to the Lebesgue measure.” Additionally, we denote by $L_{\mathbb{R}}^2(E) \subset L_{\mathbb{C}}^2(E)$ the subspace of real-valued functions. For any $F \in L_{\mathbb{C}}^2(\mathbb{R}^2)$, \bar{F} denotes its flipped version: $\bar{F}(\mathbf{x}) := F(-\mathbf{x})$.

The 2D Fourier transform of any $F \in L_{\mathbb{C}}^2(\mathbb{R}^2)$ is denoted by $\hat{F} \in L_{\mathbb{C}}^2(\mathbb{R}^2)$, such that

$$\forall \boldsymbol{\nu} \in \mathbb{R}^2, \hat{F}(\boldsymbol{\nu}) := \iint_{\mathbb{R}^2} F(\mathbf{x}) e^{-i\langle \boldsymbol{\nu}, \mathbf{x} \rangle} d^2\mathbf{x}. \quad (7)$$

For any $\varepsilon > 0$ and $\boldsymbol{\nu} \in \mathbb{R}^2$, we denote by $\mathcal{V}(\boldsymbol{\nu}, \varepsilon) \subset L_{\mathbb{C}}^2(\mathbb{R}^2)$ the set of functions whose Fourier transform is supported in a square region of size $\varepsilon \times \varepsilon$ centered in $\boldsymbol{\nu}$:

$$\mathcal{V}(\boldsymbol{\nu}, \varepsilon) := \left\{ \Psi \in L_{\mathbb{C}}^2(\mathbb{R}^2) \mid \text{supp } \hat{\Psi} \subset B_{\infty}(\boldsymbol{\nu}, \varepsilon/2) \right\}. \quad (8)$$

$\boldsymbol{\nu}$ and ε are respectively referred to as *characteristic frequency* and *bandwidth*. Finally, for any $\mathbf{h} \in \mathbb{R}^2$, we consider the translation operator, denoted by $\mathcal{T}_{\mathbf{h}}$, defined by

$$\mathcal{T}_{\mathbf{h}}F : \mathbf{x} \mapsto F(\mathbf{x} - \mathbf{h}). \quad (9)$$

Discrete Framework We denote by $l_{\mathbb{C}}^2(\mathbb{Z}^2)$ the space of 2D complex-valued square-summable sequences, represented by straight capital letters. Indexing is made between square brackets: $\forall X \in l_{\mathbb{C}}^2(\mathbb{Z}^2), \forall \mathbf{n} \in \mathbb{Z}^2, X[\mathbf{n}] \in \mathbb{C}$, and we denote by $l_{\mathbb{R}}^2(\mathbb{Z}^2) \subset l_{\mathbb{C}}^2(\mathbb{Z}^2)$ the subset of real-valued sequences. For any $V \in l_{\mathbb{C}}^2(\mathbb{Z}^2)$, \bar{V} denotes its “flipped” version as defined in (2). The convolution and subsampling operators, respectively denoted by $*$ and \downarrow , are defined in (3). 2D images, feature maps and convolution kernels are considered as elements of $l_{\mathbb{C}}^2(\mathbb{Z}^2)$. Additionally, multichannel arrays of 2D sequences are denoted by bold straight capital letters, for instance: $\mathbf{X} := (X_k)_{k \in \{0, \dots, K-1\}}$. Note that indexing starts at 0 to comply with practical implementations.

The 2D discrete-time Fourier transform of any $X \in l_{\mathbb{C}}^2(\mathbb{Z}^2)$, denoted by $\hat{X} \in L_{\mathbb{C}}^2([-\pi, \pi]^2)$, is defined by

$$\forall \boldsymbol{\theta} \in [-\pi, \pi]^2, \hat{X}(\boldsymbol{\theta}) := \sum_{\mathbf{n} \in \mathbb{Z}^2} X[\mathbf{n}] e^{-i\langle \boldsymbol{\theta}, \mathbf{n} \rangle}. \quad (10)$$

For any $\kappa \in]0, 2\pi]$ and $\boldsymbol{\theta} \in B_{\infty}(\pi)$, we denote by $\mathcal{J}(\boldsymbol{\theta}, \kappa) \subset l_{\mathbb{C}}^2(\mathbb{Z}^2)$ the set of 2D sequences whose Fourier transform is supported in a square region of size $\kappa \times \kappa$ centered in $\boldsymbol{\theta}$:

$$\mathcal{J}(\boldsymbol{\theta}, \kappa) := \left\{ W \in l_{\mathbb{C}}^2(\mathbb{Z}^2) \mid \text{supp } \hat{W} \subset B_{\infty}(\boldsymbol{\theta}, \kappa/2) \right\}. \quad (11)$$

As in the continuous framework, $\boldsymbol{\theta}$ and κ are respectively referred to as *characteristic frequency* and *bandwidth*. The elements of $\mathcal{J}(\boldsymbol{\theta}, \kappa)$ are designated as *Gabor-like filters*.

Remark 2. The support $B_{\infty}(\boldsymbol{\theta}, \kappa/2)$ actually lives in the quotient space $[-\pi, \pi]^2 / (2\pi\mathbb{Z}^2)$. Consequently, when $\boldsymbol{\theta}$ is close to an edge, a fraction of this region is located at the far end of the frequency domain. From now on, the choice of $\boldsymbol{\theta}$ and κ is implicitly assumed to avoid such a situation.




Figure 1. Spatial (left) and Fourier (right) representations of convolution kernels in the first layer of AlexNet, after training with ImageNet ILSVRC 2012-2017 (Russakovsky et al., 2015). Each kernel connects the 3 RGB input channels to one of the 64 output channels.




Figure 2. (a), (b): Real and imaginary parts of a Gabor-like filter W as defined in (12). (c), (d): Magnitude spectra (modulus of the Fourier transform) of V and W , respectively.

2.2 Intuition

In many CNNs for computer vision, input images are first transformed through subsampled (or strided) convolutions. For instance, in AlexNet, convolution kernels are of size 11×11 and the subsampling factor is equal to 4. Figure 1 displays the corresponding kernels after training with ImageNet. This linear transform is generally followed by rectified linear unit (ReLU) and max pooling.

We can observe that many kernels display oscillating patterns with well-defined orientations (Gabor-like filters). We denote by $V \in l_{\mathbb{R}}^2(\mathbb{Z}^2)$ one of these “well-behaved” filters. Its Fourier spectrum roughly consists in two bright spots which are symmetric with respect to the origin.¹ Next, we consider a complex-valued companion $W \in l_{\mathbb{C}}^2(\mathbb{Z}^2)$ such that

$$\widehat{W}(\boldsymbol{\omega}) := (1 + \operatorname{sgn}\langle \boldsymbol{\omega}, \mathbf{u} \rangle) \cdot \widehat{V}(\boldsymbol{\omega}) \quad \forall \boldsymbol{\omega} \in [-\pi, \pi]^2, \quad (12)$$

where \mathbf{u} denotes a unit vector orthogonal to the filter’s orientation.

We can show that V is the real part of W , and that $W = V + i\mathcal{H}(V)$, where \mathcal{H} denotes the two-dimensional Hilbert transform as introduced by Havlicek et al. (1997). It satisfies

$$\widehat{\mathcal{H}(V)}(\boldsymbol{\omega}) := -i \operatorname{sgn}\langle \boldsymbol{\omega}, \mathbf{u} \rangle \widehat{V}(\boldsymbol{\omega}). \quad (13)$$

As a consequence, \widehat{W} is equal to $2\widehat{V}$ on one half of the Fourier domain, and 0 on the other half. Therefore, only one bright spot remains in the spectrum. We refer the reader to fig. 2 for visual example of complex-valued Gabor-like filter. It turns out that such complex filters with high frequency resolution produce stable signal representations, as we will see in Section 2. In the subsequent sections, we then wonder whether this property is kept when considering the max pooling of real-valued convolutions.

¹Actually, the Fourier transform of any real-valued sequence is centrally symmetric: $\widehat{V}(-\boldsymbol{\omega}) = \widehat{V}(\boldsymbol{\omega})^*$. The specificity of well-oriented filters lies in the concentration of their power spectrum around two precise locations.

In what follows, W will be referred to as a discrete Gabor-like filter, and the coefficients resulting from the convolution with W will be referred to as discrete Gabor-like coefficients. The aim of this section is to show that, under the Gabor hypothesis on the convolution kernels $W \in l_{\mathbb{C}}^2(\mathbb{Z}^2)$, $\mathbb{C}\text{Mod}$ is nearly shift-invariant. To clarify, we establish that

$$U_m^{\text{mod}}[W](X) \approx U_m^{\text{mod}}[W](\mathcal{T}_{\mathbf{u}}X), \quad (14)$$

for “small” translation vectors $\mathbf{u} \in \mathbb{R}^2$, where a formal definition of the translation operator will be defined in (28). This result is hinted by Kingsbury & Magarey (1998) but not formally proven.

2.3 Continuous Framework

We introduce several results regarding functions defined on the continuous space \mathbb{R}^2 . Near-shift invariance on discrete 2D sequences will then be derived from these results by taking advantage of sampling theorems. Lemma 1 below is adapted from Waldspurger (2015, pp. 190–191).

Lemma 1. *Given $\varepsilon > 0$ and $\boldsymbol{\nu} \in \mathbb{R}^2$, let $\Psi \in \mathcal{V}(\boldsymbol{\nu}, \varepsilon)$ denote a complex-valued filter such as defined in (8). Next, for any real-valued function $F \in L_{\mathbb{R}}^2(\mathbb{R}^2)$, we consider the complex-valued function $F_0 \in L_{\mathbb{C}}^2(\mathbb{R}^2)$ defined by*

$$F_0 : \mathbf{x} \mapsto (F * \bar{\Psi})(\mathbf{x}) e^{i\langle \boldsymbol{\nu}, \mathbf{x} \rangle}. \quad (15)$$

Then F_0 is low-frequency. Specifically,

$$\text{supp } \widehat{F_0} \subset B_{\infty}(\varepsilon/2). \quad (16)$$

Proof. See Appendix A.1. □

On the other hand, the following proposition provides a shift invariance bound for low-frequency functions such as introduced above.

Proposition 1. *For any $F_0 \in L_{\mathbb{R}}^2(\mathbb{R}^2)$ such that $\text{supp } \widehat{F_0} \subset B_{\infty}(\varepsilon/2)$, and any $\mathbf{h} \in \mathbb{R}^2$,*

$$\|\mathcal{T}_{\mathbf{h}}F_0 - F_0\|_{L^2} \leq \alpha(\varepsilon\mathbf{h}) \|F_0\|_{L^2}, \quad (17)$$

where we have defined

$$\alpha : \boldsymbol{\tau} \mapsto \frac{\|\boldsymbol{\tau}\|_1}{2}. \quad (18)$$

Proof. See Appendix A.2. □

2.4 Adaptation to Discrete 2D Sequences

Given $\kappa \in]0, 2\pi]$ and $\boldsymbol{\theta} \in B_{\infty}(\pi)$, let $W \in \mathcal{J}(\boldsymbol{\theta}, \kappa)$ denote a discrete Gabor-like filter such as defined in (11). For any image $X \in l_{\mathbb{C}}^2(\mathbb{Z}^2)$ with finite support and any subsampling factor $m \in \mathbb{N} \setminus \{0\}$, we express $(X * \bar{W}) \downarrow m$ using the continuous framework introduced above, and derive an invariance formula.

For any sampling interval $s \in \mathbb{R}_{>0}$, let $\Phi^{(s)} \in L_{\mathbb{R}}^2(\mathbb{R}^2)$ denote the Shannon scaling function parameterized by s , such that

$$\widehat{\Phi^{(s)}} := s \mathbb{1}_{B_{\infty}(\pi/s)}. \quad (19)$$

This 2D function is a tensor product of scaled and normalized sinc functions. For any $\mathbf{n} \in \mathbb{Z}^2$, we denote by $\Phi_{\mathbf{n}}^{(s)}$ a shifted version of $\Phi^{(s)}$, satisfying

$$\Phi_{\mathbf{n}}^{(s)}(\mathbf{x}) := \Phi^{(s)}(\mathbf{x} - s\mathbf{n}). \quad (20)$$

Then, $\{\Phi_{\mathbf{n}}^{(s)}\}_{\mathbf{n} \in \mathbb{Z}^2}$ is an orthonormal basis of

$$\mathcal{V}^{(s)} := \{F \in L_{\mathbb{C}}^2(\mathbb{R}^2) \mid \text{supp } \widehat{F} \subset B_{\infty}(\pi/s)\}. \quad (21)$$

Then, using the notation introduced in (8), we have $\mathcal{V}^{(s)} = \mathcal{V}(\mathbf{0}, 2\pi/s)$.

We now consider the following lemma.

Lemma 2. Let $s > 0$. For any $F \in \mathcal{V}^{(s)}$ and any $\xi \in B_\infty(\pi/s)$, we have

$$\widehat{F}(\xi) = s \widehat{X}(s\xi), \quad (22)$$

where $X \in l_{\mathbb{C}}^2(\mathbb{Z}^2)$ is a uniform sampling of F , defined such that $X[\mathbf{n}] := sF(s\mathbf{n})$, for any $\mathbf{n} \in \mathbb{Z}^2$. Moreover, we have the following norm equality:

$$\|F\|_{L^2} = \|X\|_2. \quad (23)$$

Proof. See Appendix A.3. □

We then get the following proposition, which draws a bond between the discrete and continuous frameworks.

Proposition 2. Let $X \in l_{\mathbb{R}}^2(\mathbb{Z}^2)$ denote an input image with finite support, and $W \in \mathcal{J}(\theta, \kappa)$. Considering a sampling interval $s \in \mathbb{R}_{>0}$, we define F_X and $\Psi_W \in \mathcal{V}^{(s)}$ such that

$$F_X := \sum_{\mathbf{n} \in \mathbb{Z}^2} X[\mathbf{n}] \phi_{\mathbf{n}}^{(s)} \quad \text{and} \quad \Psi_W := \sum_{\mathbf{n} \in \mathbb{Z}^2} W[\mathbf{n}] \phi_{\mathbf{n}}^{(s)}. \quad (24)$$

Then,

$$\Psi_W \in \mathcal{V}(\theta/s, \kappa/s). \quad (25)$$

Moreover, for all $\mathbf{n} \in \mathbb{Z}$,

$$X[\mathbf{n}] = sF_X(s\mathbf{n}); \quad W[\mathbf{n}] = s\Psi_W(s\mathbf{n}), \quad (26)$$

and, for a given subsampling factor $m \in \mathbb{N} \setminus \{0\}$,

$$((X * \overline{W}) \downarrow m)[\mathbf{n}] = (F_X * \overline{\Psi}_W)(m\mathbf{n}). \quad (27)$$

Proof. See Appendix A.4. □

Proposition 2 introduces a latent subspace of $L_{\mathbb{R}}^2(\mathbb{R}^2)$ from which input images are uniformly sampled. This allows us to define, for any $\mathbf{u} \in \mathbb{R}^2$, a translation operator $\mathcal{T}_{\mathbf{u}}$ on discrete sequences, even if \mathbf{u} has non-integer values:

$$\mathcal{T}_{\mathbf{u}}X[\mathbf{n}] := s\mathcal{T}_{s\mathbf{u}}F_X(s\mathbf{n}), \quad (28)$$

where F_X is defined in (24). We can indeed show that this definition is independent from the choice of sampling interval $s > 0$. Moreover, given $X \in l_{\mathbb{R}}^2(\mathbb{Z}^2)$, we have

$$\forall \mathbf{p} \in \mathbb{Z}^2, \mathcal{T}_{\mathbf{p}}X[\mathbf{n}] = X[\mathbf{n} - \mathbf{p}]; \quad (29)$$

$$\forall \mathbf{u}, \mathbf{v} \in \mathbb{R}^2, \mathcal{T}_{\mathbf{u}}(\mathcal{T}_{\mathbf{v}}X) = \mathcal{T}_{\mathbf{u}+\mathbf{v}}X, \quad (30)$$

which shows that $\mathcal{T}_{\mathbf{u}}$ corresponds to the intuitive idea of a translation operator. Expressions (29) and (30) are direct consequence of the following lemma, which bonds the shift operator in the discrete and continuous frameworks.

Lemma 3. For any $X \in l_{\mathbb{R}}^2(\mathbb{Z}^2)$ and any $\mathbf{u} \in \mathbb{R}^2$,

$$F_{\mathcal{T}_{\mathbf{u}}X} = \mathcal{T}_{s\mathbf{u}}F_X. \quad (31)$$

Proof. See Appendix A.5. □

We now consider the following corollary to Proposition 2.

Corollary 1. For any shift vector $\mathbf{u} \in \mathbb{R}^2$, we have

$$((\mathcal{T}_{\mathbf{u}}X * \overline{W}) \downarrow m)[\mathbf{n}] = (\mathcal{T}_{s\mathbf{u}}F_X * \overline{\Psi}_W)(m\mathbf{n}). \quad (32)$$

Proof. Applying (27) in Proposition 2 with $X \leftarrow \mathcal{T}_{\mathbf{u}}X$, we get

$$((\mathcal{T}_{\mathbf{u}}X * \overline{W}) \downarrow m)[\mathbf{n}] = (F_{\mathcal{T}_{\mathbf{u}}X} * \overline{\Psi}_W)(m\mathbf{n}), \quad (33)$$

and Lemma 3 concludes the proof. □

2.5 Shift Invariance in the Discrete Framework

We consider the $\mathbb{C}\text{Mod}$ operator defined in (5). For the sake of conciseness, in what follows we will write U_m^{mod} instead of $U_m^{\text{mod}}[W]$, when no ambiguity is possible. First, we state the following lemma.

Lemma 4. *For any input image $X \in l_{\mathbb{R}}^2(\mathbb{Z}^2)$ with finite support, and any Gabor-like filter $W \in \mathcal{J}(\boldsymbol{\theta}, \kappa)$, we consider the low-frequency function*

$$F_0 : \mathbf{x} \mapsto (F_X * \bar{\Psi}_W)(\mathbf{x}) e^{i(\boldsymbol{\theta}/s, \mathbf{x})}, \quad (34)$$

with F_X and Ψ_W satisfying (24). If $\kappa \leq \pi/m$, then

$$F_0 \in \mathcal{V}^{(s')}. \quad (35)$$

Moreover, for any $\mathbf{h} \in \mathbb{R}^2$,

$$\sum_{\mathbf{n} \in \mathbb{Z}^2} \left| \mathcal{T}_{\mathbf{h}} F_0(s' \mathbf{n}) - F_0(s' \mathbf{n}) \right|^2 = \frac{1}{s'^2} \|\mathcal{T}_{\mathbf{h}} F_0 - F_0\|_{L^2}^2, \quad (36)$$

where we have denoted $s' := 2ms$. Finally,

$$\|U_m^{\text{mod}} X\|_2 = \frac{1}{s'} \|F_0\|_{L^2}. \quad (37)$$

Proof. See Appendix A.6. □

We are now ready to state the main result about shift invariance of $\mathbb{C}\text{Mod}$ outputs.

Theorem 1 (Shift invariance of $\mathbb{C}\text{Mod}$). *Let $W \in \mathcal{J}(\boldsymbol{\theta}, \kappa)$ denote a discrete Gabor-like filter and $m \in \mathbb{N} \setminus \{0\}$ denote a subsampling factor. Then, under the following condition:*

$$\kappa \leq \pi/m, \quad (38)$$

we have, for any input image $X \in l_{\mathbb{R}}^2(\mathbb{Z}^2)$ with finite support and any translation vector $\mathbf{u} \in \mathbb{R}^2$,

$$\|U_m^{\text{mod}}(\mathcal{T}_{\mathbf{u}} X) - U_m^{\text{mod}} X\|_2 \leq \alpha(\kappa \mathbf{u}) \|U_m^{\text{mod}} X\|_2, \quad (39)$$

where α has been defined in (18).

Proof. See Appendix A.7. □

Interestingly, the reference value used in Theorem 1, *i.e.*, $\|U_m^{\text{mod}} X\|_2$, is fully shift-invariant, as stated in the following proposition.

Proposition 3. *Let $W \in \mathcal{J}(\boldsymbol{\theta}, \kappa)$ and $m \in \mathbb{N} \setminus \{0\}$. Under condition (38), we have, for any $X \in l_{\mathbb{R}}^2(\mathbb{Z}^2)$ and any $\mathbf{u} \in \mathbb{R}^2$,*

$$\|U_m^{\text{mod}}(\mathcal{T}_{\mathbf{u}} X)\|_2 = \|U_m^{\text{mod}} X\|_2. \quad (40)$$

Proof. See Appendix A.8. □

3 From CMod to RMax

CMod operators are found in ScatterNets and complex-valued convolutional networks (Tygert et al., 2016). However, they are absent from conventional, freely-trained CNN architectures. Therefore, Theorem 1 cannot be applied as is. Instead, the first convolution layer contains real-valued kernels, and is generally followed by ReLU and max pooling. As shown in Section 5, this process can be described with RMax operators, such as defined in (1).

As explained in Section 1.1, an important number of trained convolution kernels exhibit oscillating patterns with well-defined frequencies and orientations. To elaborate, let $V \in l_{\mathbb{R}}^2(\mathbb{Z}^2)$ denote such a trained kernel, and consider $W \in l_{\mathbb{C}}^2(\mathbb{Z}^2)$ as the complex-valued companion of V satisfying (12). Then, W has its energy concentrated in a small region of the Fourier domain. We thus state the hypotheses that $W \in \mathcal{J}(\boldsymbol{\theta}, \kappa)$ (11) for a certain value of $\boldsymbol{\theta} \in [-\pi, \pi]^2$ and $\kappa \in]0, 2\pi]$. For the sake of conciseness, from now on we write $U_{m,q}^{\max}$ instead of $U_{m,q}^{\max}[W]$, when no ambiguity is possible. In what follows, we establish conditions on W under which CMod (5) and RMax (1) operators produce comparable outputs. The final goal, achieved in Section 4, is to provide a shift invariance bound for RMax.

To give an intuition about why RMax may act as a proxy for CMod, we place ourselves in the continuous framework. Consider the real-valued wavelet transform output $\text{Re } F_1 := F * \text{Re } \bar{\Psi}$, employed in RMax, as the real part of the complex-valued wavelet transform output $F_1 := F * \bar{\Psi}$, used in CMod. At a given location $\boldsymbol{x} \in \mathbb{R}^2$, the corresponding imaginary part may carry a large amount of information, which somehow needs to be retrieved. The key idea is that, if Ψ is sufficiently localized in the Fourier domain, then only the phase of F_1 significantly varies in the vicinity of \boldsymbol{x} , whereas its magnitude remains nearly constant. Therefore, finding the maximum value of $\text{Re } F_1$ within a local neighborhood around \boldsymbol{x} is nearly equivalent to shifting the phase of $F_1(\boldsymbol{x})$ towards 0. The resulting value then approximates $|F_1(\boldsymbol{x})|$. To put it differently, max pooling pushes energy towards lower frequencies, in a similar way as the modulus does for complex-valued transforms (Bruna & Mallat, 2013). This result is hinted in Section 3.1.

Regretfully, things do not work so smoothly in the discrete case. At first glance, this is surprising because Shannon’s sampling theorem allows to cast discrete problems into the continuous framework, as done in Section 2.4. However, as explained in Section 3.2, max pooling operates over a discrete grid instead of a continuous window. Consequently, in some situations, the maximum value may fall far away from any zero-phase coefficient. Taking into account this behavior, we adopt a probabilistic point of view, as detailed in Section 3.4. Then, we provide in Section 3.5 an upper bound for the expected gap between CMod and RMax outputs.

3.1 Continuous Framework

This section, inspired from Waldspurger (2015, pp. 190–191), provides an intuition about resemblance between RMax and CMod in the continuous framework. As will be highlighted in Section 3.2, adaptation to discrete 2D sequences is not straightforward and will require a probabilistic approach.

We consider an input function $F \in L_{\mathbb{R}}^2(\mathbb{R}^2)$ and a band-pass filter $\Psi \in \mathcal{V}(\boldsymbol{\nu}, \varepsilon)$. Let us also consider

$$G : (\boldsymbol{x}, \boldsymbol{h}) \mapsto \cos(\langle \boldsymbol{\nu}, \boldsymbol{h} \rangle - H(\boldsymbol{x})), \quad (41)$$

where $H : \mathbb{R}^2 \rightarrow [0, 2\pi[$ denotes the phase of $F * \bar{\Psi}$. Lemma 1 introduced low-frequency functions F_0 , with slow variations. In a nutshell, since $\text{supp } F_0 \subset B_{\infty}(\varepsilon/2)$, we can write

$$\|\boldsymbol{h}\|_2 \ll \lambda_{F_0} \implies F_0(\boldsymbol{x} + \boldsymbol{h}) \approx F_0(\boldsymbol{x}), \quad (42)$$

where we have defined $\lambda_{F_0} := 2\pi/\varepsilon$. Therefore, according to Proposition 4 below, we get the following approximation of $F * \text{Re } \bar{\Psi}$ in a neighborhood around any point $\boldsymbol{x} \in \mathbb{R}^2$:

$$\|\boldsymbol{h}\|_2 \ll \lambda_{F_0} \implies (F * \text{Re } \bar{\Psi})(\boldsymbol{x} + \boldsymbol{h}) \approx |(F * \bar{\Psi})(\boldsymbol{x})| G(\boldsymbol{x}, \boldsymbol{h}). \quad (43)$$

Proposition 4. For any $\boldsymbol{h} \in \mathbb{R}^2$,

$$|(F * \text{Re } \bar{\Psi})(\boldsymbol{x} + \boldsymbol{h}) - |(F * \bar{\Psi})(\boldsymbol{x})| G(\boldsymbol{x}, \boldsymbol{h})| \leq |F_0(\boldsymbol{x} + \boldsymbol{h}) - F_0(\boldsymbol{x})|. \quad (44)$$

Proof. See Appendix A.9. □

On the one hand, we consider a continuous equivalent of the $\mathbb{C}\text{Mod}$ operator $U_m^{\text{mod}}[\mathbb{W}]$ as introduced in (5). Such an operator, denoted by $U^{\text{mod}}[\Psi]$, is defined, for any $F \in L_{\mathbb{R}}^2(\mathbb{R}^2)$, by

$$U^{\text{mod}}[\Psi](F) : \mathbf{x} \mapsto |(F * \bar{\Psi})(\mathbf{x})|. \quad (45)$$

On the other hand, we consider the continuous counterpart of $\mathbb{R}\text{Max}$ as introduced in (1). It is defined as the maximum value of $F * \text{Re}\bar{\Psi}$ over a sliding spatial window of size $r > 0$. This is possible because F and $\text{Re}\bar{\Psi}$ both belong to $L_{\mathbb{R}}^2(\mathbb{R}^2)$, and therefore $F * \text{Re}\bar{\Psi}$ is continuous. Such an operator, denoted by $U_r^{\text{max}}[\Psi]$, is defined, for any $F \in L_{\mathbb{R}}^2(\mathbb{R}^2)$, by

$$U_r^{\text{max}}[\Psi](F) : \mathbf{x} \mapsto \max_{\|\mathbf{h}\|_{\infty} \leq r} (F * \text{Re}\bar{\Psi})(\mathbf{x} + \mathbf{h}). \quad (46)$$

For the sake of conciseness, the parameter between square brackets is ignored from now on. If $r \ll \lambda_{F_0}$, then (43) is valid for any $\mathbf{h} \in B_{\infty}(r)$. Then, using (45) and (46), we get

$$r \ll \lambda_{F_0} \implies U_r^{\text{max}}F(\mathbf{x}) \approx U^{\text{mod}}F(\mathbf{x}) \max_{\|\mathbf{h}\|_{\infty} \leq r} G(\mathbf{x}, \mathbf{h}). \quad (47)$$

Using the periodicity of G , we can show that, if $r \geq \frac{\pi}{\|\nu\|_2}$, then $\mathbf{h} \mapsto G(\mathbf{x}, \mathbf{h})$ necessarily reaches its maximum value (*i.e.*, 1) on $B_{\infty}(r)$. We therefore get

$$\frac{\pi}{\|\nu\|_2} \leq r \ll \frac{2\pi}{\varepsilon} \implies U_r^{\text{max}}F(\mathbf{x}) \approx U^{\text{mod}}F(\mathbf{x}). \quad (48)$$

3.2 Adaptation to Discrete 2D Sequences

As in Section 2.4, we consider an input image $X \in l_{\mathbb{R}}^2(\mathbb{Z}^2)$, a complex, analytic convolution kernel $W \in \mathcal{J}(\theta, \kappa)$, a subsampling factor $m \in \mathbb{N} \setminus \{0\}$ and an integer $q \in \mathbb{N} \setminus \{0\}$, referred to as a *half-size*, such that max pooling operates on a grid of size $(2q + 1) \times (2q + 1)$. We seek a relationship between

$$Y^{\text{max}} := U_{m,q}^{\text{max}}[W](X) \quad \text{and} \quad Y^{\text{mod}} := U_m^{\text{mod}}[W](X), \quad (49)$$

where $U_{m,q}^{\text{max}}[W]$ ($\mathbb{R}\text{Max}$) and $U_m^{\text{mod}}[W]$ ($\mathbb{C}\text{Mod}$) have been respectively defined in (1) and (5). As before, in what follows we omit the parameter between square brackets.

We now use the sampling results from Proposition 2. Let F_X and $\Psi_W \in \mathcal{V}^{(s)}$ denote the functions satisfying (24). Recall that the continuous versions of $\mathbb{C}\text{Mod}$ and $\mathbb{R}\text{Max}$ operators have been defined in (45) and (46), respectively. On the one hand, we apply (27) with $m \leftarrow 2m$ to Y^{mod} . For any $\mathbf{n} \in \mathbb{Z}^2$,

$$U_m^{\text{mod}}X[\mathbf{n}] = (F_X * \bar{\Psi}_W)(\mathbf{x}_{\mathbf{n}}) \quad (50)$$

$$= U^{\text{mod}}F_X(\mathbf{x}_{\mathbf{n}}), \quad (51)$$

with $\mathbf{x}_{\mathbf{n}} := 2m\mathbf{s}\mathbf{n}$. On the other hand, we postulate that

$$U_{m,q}^{\text{max}}X[\mathbf{n}] = U_r^{\text{max}}F_X(\mathbf{x}_{\mathbf{n}}) \quad (52)$$

for a certain value of $r \in \mathbb{R}_{>0}$. Then, (48) implies $Y^{\text{mod}} \approx Y^{\text{max}}$. However, as explained hereafter, (52) is not satisfied, due to the discrete nature of the max pooling grid. According to (1) and (4), we have

$$U_{m,q}^{\text{max}}X[\mathbf{n}] = \max_{\|\mathbf{p}\|_{\infty} \leq q} \text{Re}((X * \bar{W}) \downarrow m)[2\mathbf{n} + \mathbf{p}]. \quad (53)$$

Therefore, according to (27) in Proposition 2, we get

$$U_{m,q}^{\text{max}}X[\mathbf{n}] = \max_{\|\mathbf{p}\|_{\infty} \leq q} (F_X * \text{Re}\bar{\Psi}_W)(\mathbf{x}_{\mathbf{n}} + \mathbf{h}_{\mathbf{p}}), \quad (54)$$

with

$$\mathbf{x}_n := 2ms\mathbf{n} \quad \text{and} \quad \mathbf{h}_p := ms\mathbf{p}. \quad (55)$$

By considering $r_q := ms(q + \frac{1}{2})$, we get a variant of (52) in which the maximum is evaluated on a discrete grid of $(2q + 1)^2$ elements, instead of the continuous region $B_\infty(r_q)$, as defined in (46) with $r \leftarrow r_q$. As a consequence, (47) is replaced in the discrete framework by

$$q \ll 2\pi/(m\kappa) \quad \Longrightarrow \quad U_{m,q}^{\max} \mathbf{X}[\mathbf{n}] \approx U_m^{\text{mod}} \mathbf{X}[\mathbf{n}] \max_{\|\mathbf{p}\|_\infty \leq q} G_X(\mathbf{x}_n, \mathbf{h}_p), \quad (56)$$

where we have introduced, similarly to (41),

$$G_X : (\mathbf{x}, \mathbf{h}) \mapsto \cos(\langle \boldsymbol{\nu}, \mathbf{h} \rangle - H_X(\mathbf{x})), \quad (57)$$

with

$$\boldsymbol{\nu} := \boldsymbol{\theta}/s \quad \text{and} \quad H_X := \angle(F_X * \bar{\Psi}_W), \quad (58)$$

where $\angle : \mathbb{C} \rightarrow [0, 2\pi[$ denotes the phase operator. Unlike the continuous case, even if the window size r_q is large enough, the existence of $\mathbf{p} \in \{-q, \dots, q\}^2$ such that $G_X(\mathbf{x}_n, \mathbf{h}_p) = 1$ is not guaranteed, as illustrated in fig. 3 with $q = 1$. Instead, we can only seek a probabilistic estimation of the normalized mean squared error between \mathbf{Y}^{\max} and \mathbf{Y}^{mod} .

Approximation (56) implies

$$q \ll 2\pi/(m\kappa) \quad \Longrightarrow \quad \|U_m^{\text{mod}} \mathbf{X} - U_{m,q}^{\max} \mathbf{X}\|_2 \approx \|\delta_{m,q} \mathbf{X}\|_2, \quad (59)$$

where $\delta_{m,q} \mathbf{X} \in l_{\mathbb{R}}^2(\mathbb{Z}^2)$ is defined such that, for any $\mathbf{n} \in \mathbb{Z}^2$,

$$\delta_{m,q} \mathbf{X}[\mathbf{n}] := U_m^{\text{mod}} \mathbf{X}[\mathbf{n}] \left(1 - \max_{\|\mathbf{p}\|_\infty \leq q} G_X(\mathbf{x}_n, \mathbf{h}_p) \right). \quad (60)$$

Expression (59) suggests that the difference between the left and right terms can be bounded by a quantity which only depends on the product $m\kappa$ (subsampling factor \times frequency localization) and the grid half-size q . In what follows, we establish a bound characterizing this approximation, which will be provided in Proposition 5.

For the sake of conciseness, we introduce the following notations:

$$A_X : (\mathbf{x}, \mathbf{h}) \mapsto (F_X * \text{Re} \bar{\Psi}_W)(\mathbf{x} + \mathbf{h}); \quad (61)$$

$$\tilde{A}_X : (\mathbf{x}, \mathbf{h}) \mapsto |(F_X * \bar{\Psi}_W)(\mathbf{x})| G_X(\mathbf{x}, \mathbf{h}). \quad (62)$$

We now consider, for any $\mathbf{n} \in \mathbb{Z}^2$, the vectors \mathbf{h}_n^{\max} and $\mathbf{h}'_n^{\max} \in ms\{-q, \dots, q\}^2$ achieving the maximum value of $A_X(\mathbf{x}_n, \mathbf{h}_p)$ and $\tilde{A}_X(\mathbf{x}_n, \mathbf{h}_p)$ over the max pooling grid, respectively. They satisfy

$$A_X^{\max}(\mathbf{x}_n) := A_X(\mathbf{x}_n, \mathbf{h}_n^{\max}) = \max_{\|\mathbf{p}\|_\infty \leq q} A_X(\mathbf{x}_n, \mathbf{h}_p); \quad (63)$$

$$\tilde{A}_X^{\max}(\mathbf{x}_n) := \tilde{A}_X(\mathbf{x}_n, \mathbf{h}'_n^{\max}) = \max_{\|\mathbf{p}\|_\infty \leq q} \tilde{A}_X(\mathbf{x}_n, \mathbf{h}_p). \quad (64)$$

Then, according to (50) and (54), we get, for any $\mathbf{n} \in \mathbb{Z}^2$,

$$A_X^{\max}(\mathbf{x}_n) = U_{m,q}^{\max} \mathbf{X}[\mathbf{n}]; \quad (65)$$

$$\tilde{A}_X^{\max}(\mathbf{x}_n) = U_m^{\text{mod}} \mathbf{X}[\mathbf{n}] \max_{\|\mathbf{p}\|_\infty \leq q} G_X(\mathbf{x}_n, \mathbf{h}_p), \quad (66)$$

and (56) becomes

$$q \ll 2\pi/(m\kappa) \quad \Longrightarrow \quad A_X^{\max}(\mathbf{x}_n) \approx \tilde{A}_X^{\max}(\mathbf{x}_n). \quad (67)$$

Remark 3. Expression (43) implies that, if $q \ll 2\pi/(m\kappa)$, then $A_X(\mathbf{x}_n, \mathbf{h}_p) \approx \tilde{A}_X(\mathbf{x}_n, \mathbf{h}_p)$ for all $\mathbf{p} \in \{-q, \dots, q\}^2$. However, this property does not guarantee that A_X and \tilde{A}_X reach their maximum in the same exact location; *i.e.*, that $\mathbf{h}_n^{\max} = \mathbf{h}'_n^{\max}$.

The following lemma provides a bound for approximation (67).

Lemma 5. For any $\mathbf{x} \in \mathbb{R}^2$,

$$\left| A_X^{\max}(\mathbf{x}_n) - \tilde{A}_X^{\max}(\mathbf{x}_n) \right| \leq \max_{\mathbf{h} \in \{\mathbf{h}_n^{\max}, \mathbf{h}'_n^{\max}\}} \left| F_0(\mathbf{x}_n + \mathbf{h}) - F_0(\mathbf{x}_n) \right|. \quad (68)$$

Proof. See Appendix A.10. □

Before stating Proposition 5, we consider the following hypothesis:

Hypothesis 1. There exists $\mathbf{h}_0 \in \mathbb{R}^2$ with $\|\mathbf{h}_0\|_2 = \sqrt{2}qms$, such that

$$\sum_{\mathbf{n} \in \mathbb{Z}^2} \max_{\mathbf{h} \in \{\mathbf{h}_n^{\max}, \mathbf{h}'_n^{\max}\}} \left| F_0(\mathbf{x}_n + \mathbf{h}) - F_0(\mathbf{x}_n) \right|^2 \leq \sum_{\mathbf{n} \in \mathbb{Z}^2} \left| F_0(\mathbf{x}_n + \mathbf{h}_0) - F_0(\mathbf{x}_n) \right|^2. \quad (69)$$

The underlying idea is explained as follows. The absolute difference between $F_0(\mathbf{x}_n + \mathbf{h})$ and $F_0(\mathbf{x}_n)$ is more likely to increase with the norm of \mathbf{h} . For any given $\mathbf{n} \in \mathbb{Z}^2$, we have, by construction, $\|\mathbf{h}_n^{\max}\|_2 \leq \sqrt{2}qms$ and $\|\mathbf{h}'_n^{\max}\|_2 \leq \sqrt{2}qms$. Therefore, we can expect to observe

$$\max_{\mathbf{h} \in \{\mathbf{h}_n^{\max}, \mathbf{h}'_n^{\max}\}} \left| F_0(\mathbf{x}_n + \mathbf{h}) - F_0(\mathbf{x}_n) \right|^2 \leq \left| F_0(\mathbf{x}_n + \mathbf{h}_0) - F_0(\mathbf{x}_n) \right|^2. \quad (70)$$

While this might occasionally not be true, Hypothesis 1 postulates that, when summing over all the data-points, the inequality holds.

We now formally state the result characterizing approximation (59).

Proposition 5. We assume that condition (38) is satisfied: $\kappa \leq \pi/m$. Then, under Hypothesis 1,

$$\|U_m^{\text{mod}}\mathbf{X} - U_{m,q}^{\max}\mathbf{X}\|_2 \leq \|\delta_{m,q}\mathbf{X}\|_2 + \beta_q(m\kappa) \|U_m^{\text{mod}}\mathbf{X}\|_2, \quad (71)$$

where $\beta_q : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is defined by

$$\beta_q : \kappa' \mapsto q\kappa'. \quad (72)$$

Proof. See Appendix A.11. □

We now seek a probabilistic estimation of $\|\delta_{m,q}\mathbf{X}\|_2$. For this purpose, we first reformulate the problem using the unit circle $\mathbb{S}^1 \subset \mathbb{C}$, before introducing a probabilistic framework in Section 3.4.

3.3 Notations on the Unit Circle

In what follows, for any $z \in \mathbb{C} \setminus \{0\}$, we denote by $\angle z \in [0, 2\pi[$ the argument of z . For any $z, z' \in \mathbb{S}^1$, the angle between z and z' is given by $\angle(z^*z')$. We then denote by $[z, z']_{\mathbb{S}^1} \subset \mathbb{S}^1$ the arc on the unit circle going from z to z' counterclockwise:

$$[z, z']_{\mathbb{S}^1} := \{z'' \in \mathbb{S}^1 \mid \angle(z^*z'') \leq \angle(z^*z')\}. \quad (73)$$

We remind readers that \mathbf{x}_n and $\mathbf{h}_p \in \mathbb{R}^2$ have been defined in (55). By using the relation $\cos \alpha = \text{Re}(e^{i\alpha})$, (57) becomes, for any $\mathbf{n} \in \mathbb{Z}^2$ and any $\mathbf{p} \in \{-q, \dots, q\}^2$,

$$G_X(\mathbf{x}_n, \mathbf{h}_p) = \text{Re}(Z_X^*(\mathbf{x}_n) Z_p(m\boldsymbol{\theta})), \quad (74)$$




Figure 3. Search for the maximum value of $\mathbf{h} \mapsto G_X(\mathbf{x}, \mathbf{h})$ over a discrete grid of size 3×3 , *i.e.*, $q = 1$. This figure displays 3 examples with different frequencies $\nu := \theta/s$ and phases $H_X(\mathbf{x})$. Hopefully the result will be close to the true maximum (left), but there are some pathological cases in which all points in the grid fall into pits (middle and right).

where we have defined the following functions with outputs on the unit circle:

$$Z_X : \mathbf{x} \mapsto e^{iH_X(\mathbf{x})} \quad \text{and} \quad Z_p : \boldsymbol{\omega} \mapsto e^{i\langle \boldsymbol{\omega}, \mathbf{p} \rangle}, \quad (75)$$

where H_X denotes the phase of $F_X * \bar{\Psi}_W$ as introduced in (58). On the one hand, $Z_X(\mathbf{x}_n)$ is the phase (represented on the unit circle \mathbb{S}^1) of the complex wavelet transform $F_X * \bar{\Psi}_W$ at location \mathbf{x}_n . On the other hand, $Z_p(m\theta)$ approximates the phase shift between any two evaluations of $F_X * \bar{\Psi}_W$ at locations \mathbf{x}, \mathbf{x}' such that $\mathbf{x}' - \mathbf{x} = \mathbf{h}_p$. This however is only true if we assume that $\bar{\Psi}_W$ exhibits slow amplitude variations. Then, $G_X(\mathbf{x}_n, \mathbf{h}_p)$ approximates the cosine of the phase of $F_X * \bar{\Psi}_W$ at location $\mathbf{x}_n + \mathbf{h}_p$.

According to (56), $\max_{\|\mathbf{p}\|_\infty \leq q} G_X(\mathbf{x}_n, \mathbf{h}_p)$ approximates the ratio between $\mathbb{R}\text{Max}$ and $\mathbb{C}\text{Mod}$ outputs at discrete location $\mathbf{n} \in \mathbb{Z}^2$. The intuition behind this is that max pooling seeks a point in a discrete grid around \mathbf{x}_n where the phase of $F_X * \bar{\Psi}_W$ is the closest to 1, thereby maximizing the amount of energy on the real part of the signal. Assuming slow amplitude variations of $\bar{\Psi}_W$, the result therefore approximates the modulus of the complex coefficients.

To get an estimation of $\delta_{m,q}X[\mathbf{n}]$ (60), we will exploit the following property. If the phases $Z_p(m\theta)$ for $\mathbf{p} \in \{-q, \dots, q\}^2$ are well distributed on the unit circle, then the values of $G_X(\mathbf{x}_n, \mathbf{h}_p)$ are evenly spread out on $[-1, 1]$. Therefore, its maximum value is more likely to be close to 1, and (60) becomes

$$\delta_{m,q}X[\mathbf{n}] \ll U_m^{\text{mod}}X[\mathbf{n}] \quad \forall \mathbf{n} \in \mathbb{Z}^2. \quad (76)$$

Let $n_q := (2q + 1)^2$ denote the number of evaluation points for the max pooling operator. For any $\boldsymbol{\omega} \in \mathbb{R}^2$, we consider a sequence of values on \mathbb{S}^1 , denoted by $(Z_i^{(q)}(\boldsymbol{\omega}))_{i \in \{0, \dots, n_q - 1\}}$, obtained by sorting $\{Z_p(\boldsymbol{\omega})\}_{\mathbf{p} \in \{-q, \dots, q\}^2}$ (75) in ascending order of their arguments:

$$0 = H_0^{(q)}(\boldsymbol{\omega}) \leq \dots \leq H_{n_q - 1}^{(q)}(\boldsymbol{\omega}) < 2\pi, \quad (77)$$

where $H_i^{(q)}(\boldsymbol{\omega})$ denotes the phase of $Z_i^{(q)}(\boldsymbol{\omega})$. In addition, we close the loop with $H_{n_q}^{(q)}(\boldsymbol{\omega}) := 2\pi$ and $Z_{n_q}^{(q)}(\boldsymbol{\omega}) := 1$. Then, we split \mathbb{S}^1 into n_q arcs delimited by $Z_i^{(q)}(\boldsymbol{\omega})$:

$$\mathfrak{A}_i^{(q)}(\boldsymbol{\omega}) := \begin{cases} [Z_i^{(q)}(\boldsymbol{\omega}), Z_{i+1}^{(q)}(\boldsymbol{\omega})]_{\mathbb{S}^1} & \text{if } H_{i+1}^{(q)}(\boldsymbol{\omega}) - H_i^{(q)}(\boldsymbol{\omega}) < 2\pi; \\ \mathbb{S}^1 & \text{otherwise.} \end{cases} \quad (78)$$

Finally, for any $i \in \{0, \dots, n_q - 1\}$, we denote by

$$\delta H_i^{(q)} : \boldsymbol{\omega} \mapsto H_{i+1}^{(q)}(\boldsymbol{\omega}) - H_i^{(q)}(\boldsymbol{\omega}) \quad (79)$$

the function computing the angular measure of arc $\mathfrak{A}_i^{(q)}(\boldsymbol{\omega})$, for any $\boldsymbol{\omega} \in \mathbb{R}^2$.

3.4 Probabilistic Framework

From now on, input X is considered as a discrete 2D stochastic process. In order to “randomize” F_X introduced in (24), we define a continuous stochastic process from X , denoted by F_X , such that

$$\forall \mathbf{x} \in \mathbb{R}^2, F_X(\mathbf{x}) := \sum_{\mathbf{n} \in \mathbb{Z}^2} X[\mathbf{n}] \Phi_{\mathbf{n}}^{(s)}(\mathbf{x}). \quad (80)$$

Next, we consider the following stochastic processes, which are parameterized by X :

$$M_X := |F_X * \bar{\Psi}_W|; \quad H_X := \angle(F_X * \bar{\Psi}_W); \quad Z_X := e^{iH_X}, \quad (81)$$

and, for any $\mathbf{p} \in \{-q, \dots, q\}^2$,

$$G_{X, \mathbf{p}} := \text{Re}(Z_X^* Z_{\mathbf{p}}(m\boldsymbol{\theta})); \quad G_X^{\max} := \max_{\|\mathbf{p}\|_{\infty} \leq q} G_{X, \mathbf{p}}, \quad (82)$$

where the deterministic function $Z_{\mathbf{p}}$ has been defined in (75).

Remark 4. By continuous extension, $H_X(\mathbf{x})$ and $Z_X(\mathbf{x})$ are uniquely defined at \mathbf{x} such that $M_X(\mathbf{x}) = 0$.

For any $\mathbf{x} \in \mathbb{R}^2$, $F_X(\mathbf{x})$ (24) and $H_X(\mathbf{x})$ (58) are respectively drawn from $F_X(\mathbf{x})$ and $H_X(\mathbf{x})$. Then, $Z_X(\mathbf{x})$ (75) is a realization of $Z_X(\mathbf{x})$. Consequently, according to (74), $G_X(\mathbf{x}, \mathbf{h}_{\mathbf{p}})$ is a realization of $G_{X, \mathbf{p}}(\mathbf{x})$. Furthermore, according to the definition of $\mathbb{C}\text{Mod}$ in (5) and \mathbf{x}_n in (55), Proposition 2 with $m \leftarrow 2m$ implies that

$$M_X(\mathbf{x}_n) = U_m^{\text{mod}} X[\mathbf{n}]. \quad (83)$$

We remind that $\boldsymbol{\theta} \in [-\pi, \pi]^2$ and $\kappa \in]0, 2\pi]$ respectively denote the center and size of the Fourier support of the complex kernel $W \in \mathcal{J}(\boldsymbol{\theta}, \kappa)$. To compute the expected discrepancy between Y^{\max} and Y^{mod} such as introduced in (49), we assume that

$$\|\boldsymbol{\theta}\|_2 \gg 2\pi/N; \quad (84)$$

$$\|\boldsymbol{\theta}\|_2 \gg \kappa, \quad (85)$$

where $N \in \mathbb{N} \setminus \{0\}$ denotes the support size of input images. These assumptions exclude low-frequency filters from the scope of our study. We then state the following hypotheses, for which a justification is provided in Appendix B.

Hypothesis 2. For any $\mathbf{x} \in \mathbb{R}^2$, $Z_X(\mathbf{x})$ is uniformly distributed on \mathbb{S}^1 .

Hypothesis 3. For any $n \in \mathbb{N} \setminus \{0\}$ and $\mathbf{x}, \mathbf{y}_0, \dots, \mathbf{y}_{n-1} \in \mathbb{R}^2$, the random variables $M_X(\mathbf{y}_i)$ for $i \in \{0, \dots, n-1\}$ are jointly independent of $Z_X(\mathbf{x})$.

3.5 Expected Quadratic Error between $\mathbb{R}\text{Max}$ and $\mathbb{C}\text{Mod}$

In this section, we propose to estimate the expected value of the stochastic quadratic error \tilde{P}_X^2 , defined such that

$$\tilde{P}_X := \|U_m^{\text{mod}} X - U_{m,q}^{\max} X\|_2 / \|U_m^{\text{mod}} X\|_2. \quad (86)$$

According to (49), this is an estimation of the relative error between Y^{mod} and Y^{\max} .

First, let us reformulate $\delta_{m,q} X$, introduced in (60), using the probabilistic framework. According to (74) and (82), we have, for any $\mathbf{n} \in \mathbb{Z}^2$,

$$\delta_{m,q} X[\mathbf{n}] := U_m^{\text{mod}} X[\mathbf{n}] (1 - G_X^{\max}(\mathbf{x}_n)). \quad (87)$$

We now consider the stochastic process

$$Q_X := 1 - G_X^{\max}, \quad (88)$$

and the random variable

$$\tilde{Q}_X := \|\delta_{m,q} X\|_2 / \|U_m^{\text{mod}} X\|_2. \quad (89)$$

The next steps are as follows: (1) at the pixel level, show that $\mathbb{E}[\mathbf{Q}_X(\mathbf{x})^2]$ depends on the subsampling factor m and the filter frequency $\boldsymbol{\theta}$, and remains close to zero with some exceptions; (2) at the image level, show that the expected value of $\tilde{\mathbf{Q}}_X^2$ is equal to the latter quantity; (3) use Proposition 5, which implies that $\tilde{\mathbf{P}}_X \approx \tilde{\mathbf{Q}}_X$, to deduce an upper bound on the expected value of $\tilde{\mathbf{P}}_X^2$.

The first point, which is established in Proposition 6 below, is a key result of this paper. It will be used to prove Theorem 2, which corresponds to the two remaining points.

Proposition 6. *Assuming Hypothesis 2, the expected value of $\mathbf{Q}_X(\mathbf{x})^2$ is independent from the choice of $\mathbf{x} \in \mathbb{R}^2$, and*

$$\mathbb{E}[\mathbf{Q}_X(\mathbf{x})^2] = \gamma_q(m\boldsymbol{\theta})^2, \quad (90)$$

where we have defined

$$\gamma_q : \boldsymbol{\omega} \mapsto \sqrt{\frac{3}{2} + \frac{1}{4\pi} \sum_{i=0}^{n_q-1} \left(\sin \delta H_i^{(q)}(\boldsymbol{\omega}) - 8 \sin \frac{\delta H_i^{(q)}(\boldsymbol{\omega})}{2} \right)}, \quad (91)$$

with $\delta H_i^{(q)}(\boldsymbol{\omega}) \in [0, 2\pi]$ (79) being the length of arc $\mathfrak{A}_i^{(q)}(\boldsymbol{\omega})$.

Proof. For the sake of readability, in this proof we omit the argument of functions $Z_{\mathbf{p}}$ (75), $Z_i^{(q)}$, $H_i^{(q)}$ (77), $\mathfrak{A}_i^{(q)}$ (78), and $\delta H_i^{(q)}$ (79); we assume they are evaluated at $\boldsymbol{\omega} \leftarrow m\boldsymbol{\theta}$. We consider the ‘‘Lebesgue’’ Borel σ -algebra on \mathbb{S}^1 generated by $\{[z, z']_{\mathbb{S}^1} \mid z, z' \in \mathbb{S}^1\} \cup \{\mathbb{S}^1\}$, on which we have defined the angular measure ϑ such that $\vartheta(\mathbb{S}^1) := 2\pi$, and

$$\forall z, z' \in \mathbb{S}^1, \vartheta([z, z']_{\mathbb{S}^1}) := \angle(z^* z'). \quad (92)$$

For any $p \in \mathbb{N} \setminus \{0\}$, we compute the p -th moment of $\mathbf{G}_X^{\max}(\mathbf{x})$ defined in (82). By considering

$$\begin{aligned} g_{\max} : \mathbb{S}^1 &\rightarrow [-1, 1] \\ z &\mapsto \max_{\|\mathbf{p}\|_{\infty} \leq q} \operatorname{Re}(z^* Z_{\mathbf{p}}), \end{aligned} \quad (93)$$

we get $\mathbf{G}_X^{\max}(\mathbf{x}) = g_{\max}(Z_X(\mathbf{x}))$. A visual representation of g_{\max} is provided in fig. 4, for two different values of $\boldsymbol{\theta}$. According to Hypothesis 2, $Z_X(\mathbf{x})$ follows a uniform distribution on \mathbb{S}^1 . Therefore,

$$\mathbb{E}[\mathbf{G}_X^{\max}(\mathbf{x})^p] = \frac{1}{2\pi} \int_{\mathbb{S}^1} g_{\max}(z)^p d\vartheta(z), \quad (94)$$

which proves that $\mathbb{E}[\mathbf{G}_X^{\max}(\mathbf{x})^p]$ does not depend on \mathbf{x} . Let us split the unit circle \mathbb{S}^1 into the arcs $\mathfrak{A}_0^{(q)}, \dots, \mathfrak{A}_{n_q-1}^{(q)}$ such as introduced in (78):

$$\mathbb{E}[\mathbf{G}_X^{\max}(\mathbf{x})^p] = \frac{1}{2\pi} \sum_{i=0}^{n_q-1} \int_{\mathfrak{A}_i^{(q)}} g_{\max}(z)^p d\vartheta(z). \quad (95)$$

Let $i \in \{0, \dots, n_q - 1\}$. We show that

$$\forall z \in \mathfrak{A}_i^{(q)}, g_{\max}(z) = \max\left(\operatorname{Re}(z^* Z_i^{(q)}), \operatorname{Re}(z^* Z_{i+1}^{(q)})\right). \quad (96)$$

Let $z \in \mathfrak{A}_i^{(q)}$ and $i' \notin \{i, i+1\}$. We prove that

$$\operatorname{Re}(z^* Z_{i'}^{(q)}) \leq \operatorname{Re}(z^* Z_i^{(q)}) \quad \text{or} \quad \operatorname{Re}(z^* Z_{i'}^{(q)}) \leq \operatorname{Re}(z^* Z_{i+1}^{(q)}). \quad (97)$$

On the one hand, we assume that $\angle(z^* Z_{i'}^{(q)}) \leq \pi$. By design of $(Z_i^{(q)})_{i \in \{0, \dots, n_q-1\}}$, we have

$$Z_{i+1}^{(q)} \in [z, Z_{i'}^{(q)}]_{\mathbb{S}^1}. \quad (98)$$

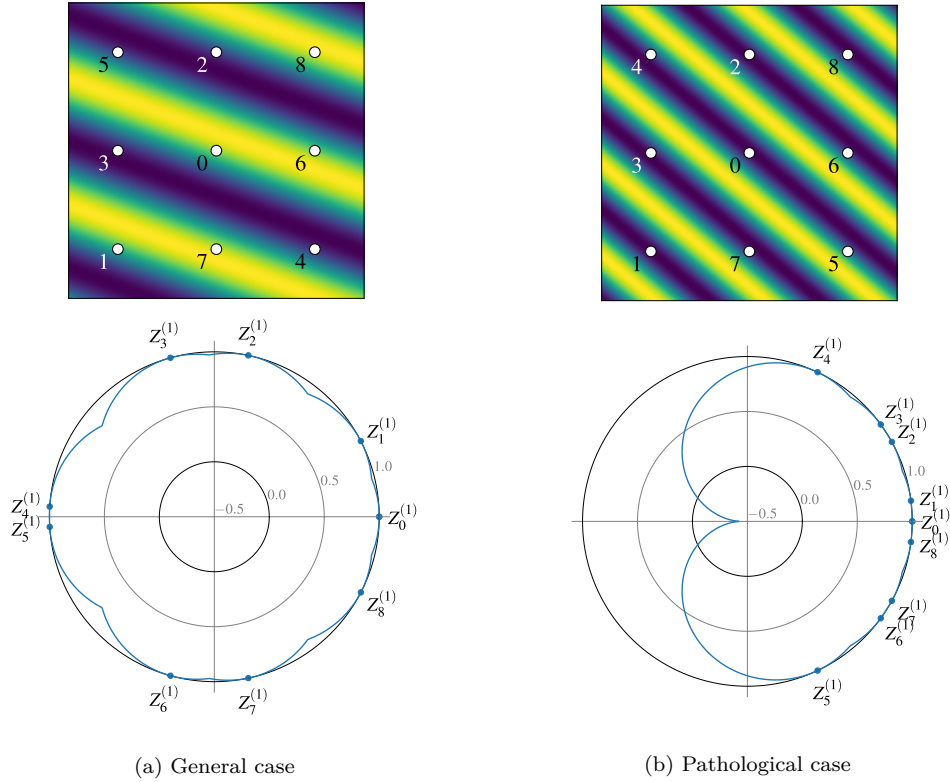


Figure 4. Top: 2D representation of $\mathbf{h} \mapsto G_X(\mathbf{x}_n, \mathbf{h})$ (57), for two different values of $\theta \in \mathbb{R}^2$, $q = 1$ and arbitrary values of $m \in \mathbb{N} \setminus \{0\}$ and $s \in \mathbb{R} \setminus \{0\}$. Assuming the plots are centered around $\mathbf{h} = \mathbf{0}$, each point materializes a location \mathbf{h}_p in the max pooling grid, for $p \in \{-q, \dots, q\}^2$. The desirable situation occurs when one of these locations falls near a ridge (bright areas), in which case the outputs produced by $\mathbb{R}\text{Max}$ and $\mathbb{C}\text{Mod}$ are similar—see (56). Each number $i \in \{0, \dots, 8\}$ represents the rank of $Z_p \in \mathbb{S}^1$ (75), when these values are sorted by ascending order of their arguments (77). If location \mathbf{h}_p gets ranked i , then we have $Z_p = Z_i^{(q)}$. Bottom: polar representations of $g_{\max} : \mathbb{S}^1 \rightarrow [-1, 1]$ (93), corresponding to the same settings. The closer the curve is from the outer ring, the more likely some points \mathbf{h}_p will fall near a ridge of G_X . (a) Case where the values Z_p are roughly evenly distributed on \mathbb{S}^1 . (b) Case where these values are concentrated in a small portion of the unit circle. The most extreme cases occurs when $Z_p = 1$ for any p . Figure 3 (middle and right) depicts two such situations.

Therefore, by definition of arcs on the unit circle (73), we get

$$\angle(z^* Z_{i+1}^{(q)}) \leq \angle(z^* Z_{i'}^{(q)}). \quad (99)$$

Then, since \cos is non-increasing on $[0, \pi]$, we get

$$\cos \angle(z^* Z_{i+1}^{(q)}) \geq \cos \angle(z^* Z_{i'}^{(q)}), \quad (100)$$

which yields the right part of (97). On the other hand, if $\angle(z^* Z_{i'}^{(q)}) \geq \pi$, a similar reasoning yields the left part of (97). Then, (96) holds.

Next, we show that, as observed in fig. 4, g_{\max} is piecewise-symmetric with respect to the center value of each arc $\mathfrak{A}_i^{(q)}$, denoted by

$$\bar{Z}_i^{(q)} := \sqrt{Z_i^{(q)} Z_{i+1}^{(q)}}. \quad (101)$$

Let $z_1, z_2 \in \mathfrak{A}_i^{(q)}$ which are symmetric with respect to $\bar{Z}_i^{(q)}$. Therefore, there exists $z' \in \mathbb{S}^1$ such that $z_1 = \bar{Z}_i^{(q)} z'$ and $z_2 = \bar{Z}_i^{(q)} z'^*$. We now prove that

$$g_{\max}(z_1) = g_{\max}(z_2). \quad (102)$$

A simple calculation yields

$$z_1^* Z_{i+1}^{(q)} = z_1'^* \tilde{Z}_i^{(q)} \quad \text{and} \quad z_2^* Z_i^{(q)} = (z_1'^* \tilde{Z}_i^{(q)})^*, \quad (103)$$

with

$$\tilde{Z}_i^{(q)} := (Z_i^{(q)*} \bar{Z}_i^{(q)}) = (\bar{Z}_i^{(q)*} Z_{i+1}^{(q)}). \quad (104)$$

Therefore,

$$\operatorname{Re}(z_1^* Z_{i+1}^{(q)}) = \operatorname{Re}(z_2^* Z_i^{(q)}). \quad (105)$$

Since z_1, z_2 both belong to $\mathfrak{A}_i^{(q)}$, $g_{\max}(z_1)$ and $g_{\max}(z_2)$ satisfy (96). Then, by symmetry, (105) implies (102). One can observe from fig. 4 that g_{\max} reaches its local minimum at the center of arc $\mathfrak{A}_i^{(q)}$, *i.e.*, $\bar{Z}_i^{(q)}$. This corresponds to a point where g_{\max} is non-differentiable.

We denote by $\bar{\mathfrak{A}}_i^{(q)} := [Z_i^{(q)}, \bar{Z}_i^{(q)}]_{\mathbb{S}^1}$ the first half of arc $\mathfrak{A}_i^{(q)}$. Then,

$$\forall z \in \bar{\mathfrak{A}}_i^{(q)}, \quad g_{\max}(z) = \operatorname{Re}(z^* Z_i^{(q)}). \quad (106)$$

As a consequence, using symmetry, we get

$$\begin{aligned} \int_{\mathfrak{A}_i^{(q)}} g_{\max}(z)^p \, d\vartheta(z) &= 2 \int_{\bar{\mathfrak{A}}_i^{(q)}} g_{\max}(z)^p \, d\vartheta(z) \\ &= 2 \int_{\bar{\mathfrak{A}}_i^{(q)}} \operatorname{Re}(z^* Z_i^{(q)})^p \, d\vartheta(z). \end{aligned}$$

By using the change of variable formula (Athreya & Lahiri, 2006, p. 81) with $z \leftarrow e^{i\eta}$, we get

$$\int_{\mathfrak{A}_i^{(q)}} g_{\max}(z)^p \, d\vartheta(z) = 2 \int_{H_i^{(q)}}^{\bar{H}_i^{(q)}} \cos^p(\eta - H_i^{(q)}) \, d\eta, \quad (107)$$

where $\bar{H}_i^{(q)} := (H_i^{(q)} + H_{i+1}^{(q)})/2$ denotes the argument of $\bar{Z}_i^{(q)}$. Then, the change of variable $\eta' \leftarrow \eta - H_i^{(q)}$ yields

$$\int_{\mathfrak{A}_i^{(q)}} g_{\max}(z)^p \, d\vartheta(z) = 2 \int_0^{\delta H_i^{(q)}/2} \cos^p \eta' \, d\eta'. \quad (108)$$

Next, we insert (108) into (95), and compute $\mathbb{E}[\mathbf{G}_X^{\max}(\mathbf{x})^p]$ for $p \leftarrow 1$ and $p \leftarrow 2$:

$$\begin{aligned} \mathbb{E}[\mathbf{G}_X^{\max}(\mathbf{x})] &= \frac{1}{\pi} \sum_{i=0}^{n_q-1} \sin \frac{\delta H_i^{(q)}}{2}; \\ \mathbb{E}[\mathbf{G}_X^{\max}(\mathbf{x})^2] &= \frac{1}{2} + \frac{1}{4\pi} \sum_{i=0}^{n_q-1} \sin \delta H_i^{(q)}. \end{aligned}$$

We recall that $\mathbf{Q}_X := 1 - \mathbf{G}_X^{\max}$. By linearity of \mathbb{E} , we get

$$\mathbb{E}[\mathbf{Q}_X(\mathbf{x})^2] := \frac{3}{2} + \frac{1}{4\pi} \sum_{i=0}^{n_q-1} \left(\sin \delta H_i^{(q)} - 8 \sin \frac{\delta H_i^{(q)}}{2} \right), \quad (109)$$

which concludes the proof. \square

We consider an ideal scenario where $(Z_i^{(q)}(m\boldsymbol{\theta}))_{i \in \{0, \dots, n_q-1\}}$ are evenly spaced on \mathbb{S}^1 . Then, an order 2 Taylor expansion yields

$$\gamma_q(m\boldsymbol{\theta}) = o(1/q^2), \quad (110)$$




Figure 5. $\gamma(m\boldsymbol{\theta})^2$ as a function of the kernel characteristic frequency $\boldsymbol{\theta} \in [-\pi, \pi]^2$. According to Theorem 2, this quantity provides an approximate bound for the expected quadratic error between $\mathbb{R}\text{Max}$ and $\mathbb{C}\text{Mod}$ outputs. The subsampling factor m has been set to 2 as in ResNet (left), and 4 as in AlexNet (right). The bright regions correspond to frequencies for which the two outputs are expected to be similar. However, in the dark regions, pathological cases such as illustrated in fig. 3 are more likely to occur.

providing an order-two-polynomial decay rate for $\mathbb{Q}_X(\mathbf{x})$, when the grid half-size q increases. Figure 5 displays $\boldsymbol{\theta} \mapsto \gamma_q(m\boldsymbol{\theta})^2$ for $\boldsymbol{\theta} \in [-\pi, \pi]^2$, with $m = 4$ and $q = 1$ as in AlexNet. We notice that, for the major part of the Fourier domain, γ_q remains close to 0. However, we observe a regular pattern of dark regions, which correspond to pathological frequencies where the repartition of $(Z_i^{(q)}(m\boldsymbol{\theta}))_{i \in \{0, \dots, n_q-1\}}$ is unbalanced.

So far, we established a result at the pixel level. Before stating Theorem 2, which extends the result to the image level, we need the following intermediate statement.

Proposition 7. *We consider the random variable*

$$\tilde{\mathcal{S}}_X := \|U_m^{\text{mod}} X\|_2. \quad (111)$$

Under Hypothesis 3, for any $\mathbf{x} \in \mathbb{R}^2$,

- $Z_X(\mathbf{x})$ is independent of $\tilde{\mathcal{S}}_X$;
- $Z_X(\mathbf{x}), M_X(\mathbf{x})$ are conditionally independent given $\tilde{\mathcal{S}}_X$.

Proof. See Appendix A.12. □

Finally, Propositions 6 and 7 yield the following theorem. It provides an upper bound on the expected value of the normalized mean squared error $\tilde{\mathcal{P}}_X^2$, such as defined in (86).

Theorem 2 (MSE between $\mathbb{C}\text{Mod}$ and $\mathbb{R}\text{Max}$). *Let $W \in \mathcal{J}(\boldsymbol{\theta}, \kappa)$ denote a discrete Gabor-like filter, $m \in \mathbb{N} \setminus \{0\}$ a subsampling factor and $q \in \mathbb{N} \setminus \{0\}$ a grid half-size. We consider a stochastic process X whose realizations are elements of $l_{\mathbb{R}}^2(\mathbb{Z}^2)$. We assume that condition (38) is satisfied: $\kappa \leq \pi/m$. Then, under Hypotheses 1 to 3,²*

$$\mathbb{E}[\tilde{\mathcal{P}}_X^2] \leq (\beta_q(m\kappa) + \gamma_q(m\boldsymbol{\theta}))^2, \quad (112)$$

where $\tilde{\mathcal{P}}_X^2$ (86) denotes the stochastic quadratic error between $\mathbb{C}\text{Mod}$ and $\mathbb{R}\text{Max}$ outputs. We remind that β_q and γ_q have been introduced in (72) and (91), respectively.

Proof. See Appendix A.13. □

²We can easily prove that these properties are independent from the choice of sampling interval $s > 0$.

Let us analyze the bound obtained in (112). The first term, $\beta_q(m\kappa)$, accounts for the localized property of the convolution filter W . This term decreases linearly with the product $m\kappa$. In the limit case where $\kappa = 0$ (infinite, nonlocal filter), we get $\beta_q(m\kappa) = 0$. Note that a smaller subsampling factor m allows for a larger bandwidth κ . Moreover, $\beta_q(m\kappa)$ increases linearly with the size of the max pooling grid, which is characterized by q . The second term, $\gamma_q(m\theta)$, accounts for the discrete nature of the max pooling grid. It strongly depends on the characteristic frequency θ , as illustrated in fig. 5. According to (110), this term has a polynomial decay when q increases. However, increasing the size of the max pooling grid also results in increasing the term $\beta_q(m\kappa)$, as explained above. Therefore, a tradeoff must be found to get an optimal bound.

4 Shift Invariance of $\mathbb{R}\text{Max}$ Outputs

In this section, we present the main theoretical claim of this paper. Based on the previous results, we provide a probabilistic measure of shift invariance for $\mathbb{R}\text{Max}$ operators. First, we state the following lemma.

Lemma 6. *If Hypotheses 2 and 3 are satisfied, then they are also true with $X \leftarrow \mathcal{T}_{\mathbf{u}}X$, for any $\mathbf{u} \in \mathbb{R}^2$.*

Proof. See Appendix A.14. □

We are now ready to state the main result about shift invariance of $\mathbb{R}\text{Max}$ outputs.

Theorem 3 (Shift invariance of $\mathbb{R}\text{Max}$). *We assume that the requirements stated in Theorem 2 are satisfied. Additionally, given a translation vector $\mathbf{u} \in \mathbb{R}^2$, we consider the following random variable:*

$$\tilde{\mathbf{R}}_{X, \mathbf{u}} := \|U_{m,q}^{\max}(\mathcal{T}_{\mathbf{u}}X) - U_{m,q}^{\max}X\|_2 / \|U_m^{\text{mod}}X\|_2. \quad (113)$$

Then, under condition (38), we have

$$\mathbb{E}[\tilde{\mathbf{R}}_{X, \mathbf{u}}] \leq 2(\beta_q(m\kappa) + \gamma_q(m\theta)) + \alpha(\kappa\mathbf{u}), \quad (114)$$

where α , β_q and γ_q are defined in (18), (72) and (91), respectively.

Proof. See Appendix A.15. □

In the bound established in (114), the sum $\beta_q(m\kappa) + \gamma_q(m\theta)$ accounts for the discrepancy between $\mathbb{R}\text{Max}$ and $\mathbb{C}\text{Mod}$ outputs, as stated in Theorem 2, whereas the term $\alpha(\kappa\mathbf{u})$ characterizes the stability of $\mathbb{C}\text{Mod}$ outputs, as stated in Theorem 1. If κ is sufficiently small, then $\alpha(\kappa\mathbf{u})$ and $\beta_q(m\kappa)$ become negligible with respect to $\gamma_q(m\theta)$, and the bound can be approximated by $2\gamma_q(m\theta)$. Theorem 3 therefore provides a validity domain for shift invariance of $\mathbb{R}\text{Max}$ operators, as illustrated in fig. 5 with $q = 1$.

Remark 5. The stochastic discrepancy introduced in (113) is estimated relatively to the $\mathbb{C}\text{Mod}$ output. This choice is motivated by the perfect shift invariance of its norm, as shown in Proposition 3.

Remark 6. In practice, most of the time max pooling is performed on a grid of size 3×3 ; therefore $q = 1$. For the sake of conciseness, we shall sometimes drop q in the notations, which implicitly means $q = 1$.

5 Adaptation to Multichannel Convolution Operators

In this section, we adapt Theorems 1 to 3 to multichannel inputs (*e.g.*, RGB images), employed in conventional CNNs such as AlexNet or ResNet.

First, we define multichannel $\mathbb{R}\text{Max}$ and $\mathbb{C}\text{Mod}$ operators relatively to (1) and (5). We denote by K and $L \in \mathbb{N} \setminus \{0\}$ the number of input and output channels, respectively. Additionally, we consider a *multichannel convolution tensor*

$$\mathbf{W} := (W_{lk})_{l \in \{0, \dots, L-1\}, k \in \{0, \dots, K-1\}} \in (l_{\mathbb{C}}^2(\mathbb{Z}^2))^{L \times K}. \quad (115)$$

Multichannel $\mathbb{R}\text{Max}$ and $\mathbb{C}\text{Mod}$ operators take as input images, denoted by

$$\mathbf{X} := (X_k)_{k \in \{0, \dots, K-1\}} \in (l_{\mathbb{R}}^2(\mathbb{Z}^2))^K. \quad (116)$$

They are defined, for any given output channel $l \in \{0, \dots, L-1\}$, by

$$U_{m,q,l}^{\max}[\mathbf{W}] : \mathbf{X} \mapsto \text{MaxPool}_q \left(\sum_{k=0}^{K-1} (X_k * \text{Re } \overline{W}_{lk}) \downarrow m \right); \quad (117)$$

$$U_{m,l}^{\text{mod}}[\mathbf{W}] : \mathbf{X} \mapsto \left| \sum_{k=0}^{K-1} (X_k * \overline{W}_{lk}) \downarrow (2m) \right|, \quad (118)$$

where $m, q \in \mathbb{N} \setminus \{0\}$ respectively denote a subsampling factor and the max pooling grid half-size. Analogously to (49) for single-channel inputs, we now consider

$$Y_l^{\max} := U_{m,q,l}^{\max}[\mathbf{W}](\mathbf{X}) \quad \text{and} \quad Y_l^{\text{mod}} := U_{m,l}^{\text{mod}}[\mathbf{W}](\mathbf{X}). \quad (119)$$

Again, in what follows we omit the parameter between square brackets. To apply Theorems 1 to 3 to the current setting on the l -th output channel, we need the following hypotheses.

Hypothesis 4 (Monochrome filters). Let

$$\widetilde{W}_l := \frac{1}{K} \sum_{k=0}^{K-1} W_{lk} \quad (120)$$

denote the mean kernel of the l -th output channel. Then, there exists $\boldsymbol{\mu}_l \in \mathbb{R}^K$ such that

$$\forall k \in \{0, \dots, K-1\}, W_{lk} = \mu_{lk} \widetilde{W}_l. \quad (121)$$

Hypothesis 5 (Gabor-like filters). There exists a bandwidth $\kappa > 0$ satisfying $\kappa \leq \pi/m$ and a frequency vector $\boldsymbol{\theta}_l \in [-\pi, \pi]^2$ such that

$$\widetilde{W}_l \in \mathcal{J}(\boldsymbol{\theta}_l, \kappa). \quad (122)$$

Note that the bandwidth κ is not indexed by l , because it shall later be assumed to be shared across the output channels. Then, under Hypothesis 4, Y_l^{\max} and Y_l^{mod} are the outputs of single-channel $\mathbb{R}\text{Max}$ and $\mathbb{C}\text{Mod}$ operators, as introduced in (1) and (5):

$$Y_l^{\max} = U_{m,q}^{\max}[\widetilde{W}_l](X_l^{\text{lum}}) \quad \text{and} \quad Y_l^{\text{mod}} = U_m^{\text{mod}}[\widetilde{W}_l](X_l^{\text{lum}}), \quad (123)$$

where $X_l^{\text{lum}} \in l_{\mathbb{R}}^2(\mathbb{Z}^2)$ (“luminance” image) is defined as the following linear combination:

$$X_l^{\text{lum}} := \sum_{k=0}^{K-1} \mu_{lk} X_k. \quad (124)$$

The results established for single-channel inputs can therefore be extended to multichannel operators. Specifically, we get the following corollaries to Theorems 1 to 3.

Corollary 2 (Shift invariance of $\mathbb{C}\text{Mod}$). *For a given output channel $l \in \{0, \dots, L-1\}$, we postulate Hypotheses 4 and 5. Then, for any input image $\mathbf{X} \in (l_{\mathbb{R}}^2(\mathbb{Z}^2))^K$ with finite support and any translation vector $\mathbf{u} \in \mathbb{R}^2$,*

$$\|U_{m,l}^{\text{mod}}(\mathcal{T}_{\mathbf{u}}\mathbf{X}) - U_{m,l}^{\text{mod}}\mathbf{X}\|_2 \leq \alpha(\kappa\mathbf{u}) \|U_{m,l}^{\text{mod}}\mathbf{X}\|_2, \quad (125)$$

where α has been defined in (18).

Corollary 3 (MSE between $\mathbb{C}\text{Mod}$ and $\mathbb{R}\text{Max}$). *As in Corollary 2, we postulate Hypotheses 4 and 5. Again, we assume that condition (38) is satisfied: $\kappa \leq \pi/m$. Additionally, we consider \mathbf{X} as a stack of K discrete stochastic processes, and assume Hypotheses 1 to 3 with $\mathbf{X} \leftarrow X_l^{\text{lum}}$ and $\mathbf{W} \leftarrow \widetilde{W}_l$. Then,*

$$\mathbb{E}[\widetilde{\mathbf{P}}_{\mathbf{X},l}^2] \leq (\beta_q(m\kappa) + \gamma_q(m\boldsymbol{\theta}_l))^2, \quad (126)$$

where we have defined the following random variable:

$$\widetilde{\mathbf{P}}_{\mathbf{X},l} := \|U_{m,l}^{\text{mod}}\mathbf{X} - U_{m,l}^{\max}\mathbf{X}\|_2 / \|U_{m,l}^{\text{mod}}\mathbf{X}\|_2. \quad (127)$$

Corollary 4 (Shift invariance of $\mathbb{R}\text{Max}$). *We assume that the requirements stated in Corollary 3 are satisfied. Then, for any translation vector $\mathbf{u} \in \mathbb{R}^2$,*

$$\mathbb{E}[\tilde{\mathbf{R}}_{\mathbf{X}, \mathbf{u}, l}] \leq 2 (\beta_q(m\kappa) + \gamma_q(m\boldsymbol{\theta}_l)) + \alpha(\kappa\mathbf{u}), \quad (128)$$

where we have defined the following random variable:

$$\tilde{\mathbf{R}}_{\mathbf{X}, \mathbf{u}, l} := \|U_{m,l}^{\max}(\mathcal{T}_{\mathbf{u}}\mathbf{X}) - U_{m,l}^{\max}\mathbf{X}\|_2 / \|U_{m,l}^{\text{mod}}\mathbf{X}\|_2. \quad (129)$$

Remark 7. In the above results, we used a translation operator on multichannel tensors, obtained by applying $\mathcal{T}_{\mathbf{u}}$, as defined in (28), to each channel \mathbf{X}_k .

6 A Case Study Implementing the Dual-Tree Complex Wavelet Packet Transform

In this section, we experimentally validate the results stated in Theorems 1 to 3. To this end, we consider a fully-deterministic scenario implementing the dual-tree complex wavelet packet transform (DT-CWPT), which exhibit characteristics akin to those observed in the initial convolution layer of freely-trained CNNs such as AlexNet or ResNet. In particular, as stated in Section 6.1, DT-CWPT achieves subsampled convolutions with oriented band-pass filters tiling the Fourier domain into overlapping square windows. As such, it provides a convenient framework to experimentally validate our theoretical findings in a controlled environment. Then, in Section 6.2, we build $\mathbb{C}\text{Mod}$ and $\mathbb{R}\text{Max}$ operators based on DT-CWPT convolution kernels.

6.1 Main Properties

In what follows, we outline the principal characteristics of DT-CWPT. A detailed description of the transform itself is provided in Appendix C.1, whereas the results presented hereafter are formally established in Appendices C.2 and C.3.

For a given decomposition depth $J \in \mathbb{N} \setminus \{0\}$, DT-CWPT achieves subsampled convolutions with 4×4^J oriented band-pass filters that tile the Fourier domain into overlapping square windows of size

$$\kappa_J := \pi/m_J, \quad \text{with} \quad m_J := 2^{J-1}. \quad (130)$$

More specifically, considering an input image $\mathbf{X} \in l_{\mathbb{R}}^2(\mathbb{Z}^2)$, it produces a set of 4×4^J output feature maps

$$\mathbf{D}^{(J)} := (\mathbf{D}_l^{\nearrow(J)}, \mathbf{D}_l^{\searrow(J)}, \mathbf{D}_l^{\swarrow(J)}, \mathbf{D}_l^{\nwarrow(J)})_{l \in \{0, \dots, 4^J - 1\}}, \quad (131)$$

where each arrow points to the Fourier quadrant where the feature map's energy is concentrated. Moreover, as stated in Proposition 10, for any $l \in \{0, \dots, 4^J - 1\}$, there exists $\mathbf{W}_l^{\nearrow(J)} \in l_{\mathbb{C}}^2(\mathbb{Z}^2)$ such that

$$\mathbf{D}_l^{\nearrow(J)} = (\mathbf{X} * \overline{\mathbf{W}_l^{\nearrow(J)}}) \downarrow 2^J. \quad (132)$$

An interesting property is that each kernel $\mathbf{W}_l^{\nearrow(J)}$ approximately satisfies

$$\mathbf{W}_l^{\nearrow(J)} \in \mathcal{J}(\boldsymbol{\theta}_l^{(J)}, \kappa_J) \quad (133)$$

for a certain characteristic frequency $\boldsymbol{\theta}_l^{(J)} \in [0, \pi]^2$. In other words, it approximately behaves as a Gabor-like filter in the discrete framework (11). Moreover, each kernel corresponds to a different frequency, thereby covering the top-right quadrant of the Fourier domain. Similar results can be established for the other three Fourier quadrants. Graphical representations of $\mathbf{W}^{\nearrow(J)} := (\mathbf{W}_l^{\nearrow(J)})_{l \in \{0, \dots, 4^J - 1\}}$ and $\mathbf{W}^{\searrow(J)} := (\mathbf{W}_l^{\searrow(J)})_{l \in \{0, \dots, 4^J - 1\}}$ are provided in fig. 6 with $J = 2$ (fig. 6a, 32 filters) and $J = 3$ (fig. 6b, 128 filters).

The $\mathbb{R}\text{Max}$ and $\mathbb{C}\text{Mod}$ operators implemented in our experiments respectively satisfy (1) and (5) with $\mathbf{W} \leftarrow \mathbf{W}_l^{\nearrow(J)}$ or $\mathbf{W}_l^{\searrow(J)}$, and $m \leftarrow m_J$. Note that increasing the decomposition depth J , and therefore the subsampling factor m_J , results in a decreased Fourier support size κ_J , therefore matching the condition stated in (38) $\kappa \leftarrow \kappa_J$ and $m \leftarrow m_J$.




Figure 6. Real part of the convolution kernels $\mathbf{W}^{\nearrow(J)}$, $\mathbf{W}^{\searrow(J)}$, with $J = 2$ (32 filters, $m_J = 2$) and $J = 3$ (128 filters, $m_J = 4$), respectively. The kernels have been computed using Q-shift orthogonal QMFs of length 10 (Kingsbury, 2003). The kernels have been respectively cropped to size 11×11 and 19×19 , for the sake of legibility. Note that the filters displayed in (a) and (b) share similarities with those found in, respectively, ResNet ($m = 2$) and AlexNet ($m = 4$), after training with ImageNet.

Remark 8. Because X is real-valued, the feature maps $D_l^{\swarrow(J)}$ and $D_l^{\nwarrow(J)}$ are the respective complex conjugates of $D_l^{\nearrow(J)}$ and $D_l^{\searrow(J)}$, and thus do not need to be explicitly computed. Then, we can easily show that $W_l^{\swarrow(J)}$ and $W_l^{\nwarrow(J)}$ are also the complex conjugates of $W_l^{\nearrow(J)}$ and $W_l^{\searrow(J)}$, respectively.

6.2 DT-CWPT-Based $\mathbb{R}\text{Max}$ and $\mathbb{C}\text{Mod}$ Operators

According to (130), (132), and (133), we can apply Theorems 1 to 3 to the dual-tree framework. More precisely, for any output channel $l \in \{0, \dots, 4^J - 1\}$, we consider the following $\mathbb{R}\text{Max}$ and $\mathbb{C}\text{Mod}$ operators:

$$U_l^{\max \nearrow} : X \mapsto \text{MaxPool} \left(\left(X * \text{Re} \bar{W}_l^{\nearrow(J)} \right) \downarrow 2^{J-1} \right); \quad (134)$$

$$U_l^{\text{mod} \nearrow} : X \mapsto \left| \left(X * \bar{W}_l^{\nearrow(J)} \right) \downarrow 2^J \right|. \quad (135)$$

Using the notations introduced in (5) and (1), we have

$$U_l^{\max \nearrow} = U_{m_J}^{\max} [W_l^{\nearrow(J)}] \quad \text{and} \quad U_l^{\text{mod} \nearrow} := U_{m_J}^{\text{mod}} [W_l^{\nearrow(J)}], \quad (136)$$

where we have defined $m_J := 2^{J-1}$. Note that, following Remark 6, we have omitted the grid half-size q , which is equal to 1 (max pooling operates on a grid of size 3×3). Furthermore, for the sake of brevity, we have omitted the depth J in the above notations.

Remark 9. Both $U_l^{\max \nearrow}$ and $U_l^{\text{mod} \nearrow}$ are implemented using DT-CWPT with J decomposition stages. However, in (134), the subsampling factor is equal to 2^{J-1} , instead of 2^J , as stated in (132). In order to accommodate this property of $\mathbb{R}\text{Max}$ operators, the last stage of DT-CWPT decomposition is carried

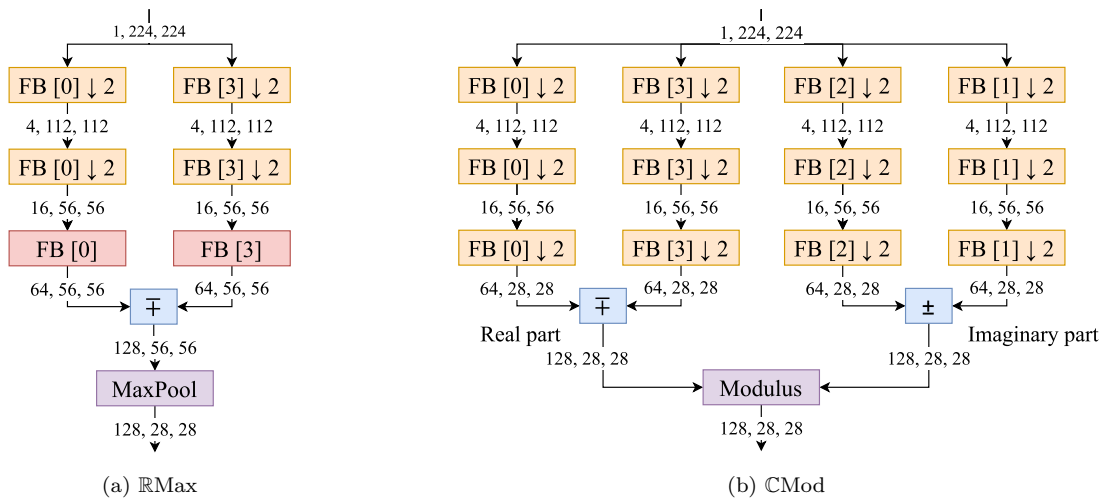


Figure 7. Detailed illustration of the $\mathbb{R}\text{Max}$ (a) and $\mathbb{C}\text{Mod}$ (b) operators based on DT-CWPT, with $J = 3$ decomposition stages. The numbers between modules correspond to the number of feature maps, height and width. The orange modules represent subsampled convolutions using one of the four 2D filter banks $\mathbf{G}^{[0-3]}$, such as introduced in (243). The FB index is indicated between square brackets. The $\mathbb{R}\text{Max}$ model (a) only computes the real part of the dual-tree coefficients, and the last stage of decomposition is performed without subsampling (red modules). Additionally, the blue modules represent linear combinations of feature maps such as described in (248).

out without subsampling, resulting in higher redundancy. This is similar to the concept of stationary wavelet transform as described by Nason & Silverman (1995). Furthermore, only the real component of the wavelet feature maps is preserved. On the other hand, $U_l^{\text{mod}\nearrow}$ implements a fully-decimated wavelet packet transform, and keeps both real and imaginary parts. Figure 7 illustrates these technical details.

6.3 Experiments and Results

We implemented the $\mathbb{R}\text{Max}$ and $\mathbb{C}\text{Mod}$ operators $U_l^{\text{max}\nearrow}$ and $U_l^{\text{mod}\nearrow}$, as introduced in (134) and (135), with both $J = 2$ and 3 stages of wavelet packet decomposition. To cover the whole frequency plane, we also implemented similar operators, denoted by $U_l^{\text{max}\searrow}$ and $U_l^{\text{mod}\searrow}$. They are associated with the convolution filters $W_l^{\searrow(J)}$, introduced in Proposition 10, with energy being located in the bottom-right quadrant. However, as explained in Remark 8, we did not need to deal with the two other quadrants (negative x -values). Using the validation set of ImageNet-1K (Russakovsky et al., 2015), ($N := 50\,000$ images), we measured the mean discrepancy between $\mathbb{R}\text{Max}$ and $\mathbb{C}\text{Mod}$ outputs, and evaluated the shift invariance of both models. Dual-tree decompositions have been performed with Q-shift orthogonal filters of length 10 (Kingsbury, 2003).

6.3.1 MSE between $\mathbb{R}\text{Max}$ and $\mathbb{C}\text{Mod}$

Each image $n \in \{0, \dots, N - 1\}$ in the dataset was converted to grayscale, from which a center crop of size 224×224 was extracted. We denote by $X_n \in l_{\mathbb{R}}^2(\mathbb{Z}^2)$ the resulting input feature map. For any $l \in \{0, \dots, 4^J - 1\}$, we denote by

$$Y_{nl}^{\text{max}\nearrow} := U_l^{\text{max}\nearrow}(X_n) \quad \text{and} \quad Y_{nl}^{\text{mod}\nearrow} := U_l^{\text{mod}\nearrow}(X_n) \quad (137)$$

the outputs of the l -th $\mathbb{R}\text{Max}$ and $\mathbb{C}\text{Mod}$ operators as defined in (134) and (135), respectively. We adopt similar notations for the bottom-right Fourier quadrant. Then, the normalized mean squared error between $Y_{nl}^{\text{mod}\nearrow}$ and $Y_{nl}^{\text{max}\nearrow}$ was computed. It is defined by the square of

$$\rho_{nl}^{\nearrow} := \|Y_{nl}^{\text{mod}\nearrow} - Y_{nl}^{\text{max}\nearrow}\|_2 / \|Y_{nl}^{\text{mod}\nearrow}\|_2. \quad (138)$$




Figure 8. Empirical estimates of the normalized mean squared error between $\mathbb{R}\text{Max}$ and $\mathbb{C}\text{Mod}$ outputs, computed on ImageNet-1K (validation set). For each channel $l \in \{0, \dots, 4^J - 1\}$, $\tilde{\rho}_l^{\nearrow 2}$ is plotted as a grayscale pixel centered in $\theta_l^{(J)}$ such as introduced in eq. (133) (top-right quadrant). Similarly, $\tilde{\rho}_l^{\searrow 2}$ is plotted in the bottom-right quadrant. Finally, the bottom- and top-left quadrants ($\tilde{\rho}_l^{\swarrow 2}$ and $\tilde{\rho}_l^{\nwarrow 2}$) are simply obtained by symmetrizing the figures. Since the subsampling factor m_J is equal to 2^{J-1} , these experimental results can be compared with the left and right parts of fig. 5. Note that the low-pass filters have been discarded because they are outside the scope of this study.

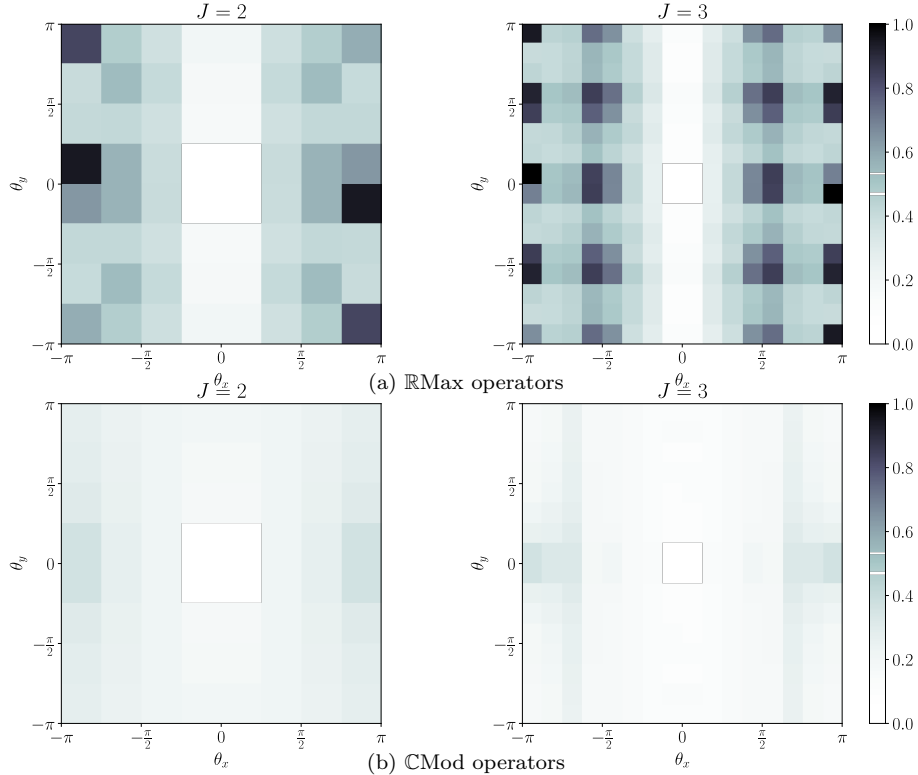


Figure 9. Shift invariance of $\mathbb{R}\text{Max}$ and $\mathbb{C}\text{Mod}$ outputs, computed on ImageNet 2012 (validation set). For each $l \in \{0, \dots, 4^J - 1\}$, $\tilde{\rho}_l^{\nearrow \max}$ (fig. 9a) and $\tilde{\rho}_l^{\nearrow \text{mod}}$ (fig. 9b) are plotted by applying the same procedure as in fig. 8.

Finally, for each output channel l , an empirical estimate for $\mathbb{E}[\tilde{\mathbb{P}}_{\mathbf{X}}^2]$, introduced in (86), was obtained by averaging $\rho_{nl}^{\nearrow 2}$ over the whole dataset. We denote by $\tilde{\rho}_l^{\nearrow 2}$ the corresponding quantity.

Since $U_l^{\max\prime\searrow}$ and $U_l^{\text{mod}\prime\searrow}$ are parameterized by $W_l^{\nearrow(J)}$, it follows that $\tilde{\rho}_l^{\nearrow 2}$ depends on the filter’s characteristic frequency $\theta_l^{(J)}$ (133). According to Proposition 11, these frequencies form a regular grid in the top-right quadrant of Fourier domain. This provides a visual representation of $\tilde{\rho}_l^{\nearrow 2}$, as shown in fig. 8. This figure also displays $\tilde{\rho}_l^{\searrow 2}$, corresponding to the bottom-right quadrant. The half-plane of negative x -values has simply been symmetrized, following Remark 8. We can observe a regular pattern of dark spots. More precisely, high discrepancies between max pooling and modulus seem to occur when the energy of $W_l^{\nearrow(J)}$ or $W_l^{\searrow(J)}$ overlaps a dark region of fig. 5. This result corroborates Theorem 3, which states that high discrepancies are expected for certain pathological frequencies, due to the search for a maximum value over a discrete grid.

6.3.2 Shift invariance

For each input image previously converted to grayscale, two crops of size 224×224 were extracted, such that the corresponding sequences X_n and X'_n are shifted by one pixel along the x -axis. From these inputs, the following quantity was then computed:

$$\rho_{nl}^{\max\prime\searrow} := \|\mathbf{Y}_{nl}^{\max\prime\searrow} - \mathbf{Y}_{nl}^{\text{mod}\prime\searrow}\|_2 / \|\mathbf{Y}_{nl}^{\text{mod}\prime\searrow}\|_2, \quad (139)$$

where $\mathbf{Y}_{nl}^{\max\prime\searrow}$ satisfies (136) with $X_n \leftarrow X'_n$. Finally, for each output channel $l \in \{0, \dots, 4^J - 1\}$, an empirical estimate for $\mathbb{E}[\tilde{\mathbb{R}}_{\mathbf{X}, \mathbf{u}}]$, satisfying (113) with $\mathbf{u} = (1, 0)^\top$, was obtained by averaging $\rho_{nl}^{\max\prime\searrow}$ over the whole dataset. We denote by $\tilde{\rho}_l^{\max\prime\searrow}$ the corresponding quantity. We point out that shift invariance is measured relatively to the norm of the CMod output, as explained in Remark 5.

On the other hand, the same procedure was applied to the CMod operators:

$$\rho_{nl}^{\text{mod}\prime\searrow} := \|\mathbf{Y}_{nl}^{\text{mod}\prime\searrow} - \mathbf{Y}_{nl}^{\text{mod}\searrow}\|_2 / \|\mathbf{Y}_{nl}^{\text{mod}\searrow}\|_2, \quad (140)$$

and $\tilde{\rho}_l^{\text{mod}\prime\searrow}$ was obtained as before by averaging $\rho_{nl}^{\text{mod}\prime\searrow}$ over the whole dataset.

A visual representation of $\tilde{\rho}_l^{\max\prime\searrow}$ and $\tilde{\rho}_l^{\text{mod}\prime\searrow}$ are provided in fig. 9 (as well as the other Fourier quadrants). Two observations can be drawn here. (1) When the filter is horizontally oriented, the corresponding output is highly stable with respect to horizontal shifts. This can be explained by noticing that such kernels perform low-pass filtering along the x -axis. The exact transposed phenomenon occurs for vertical shifts. (2) Elsewhere, we observe that high discrepancies between RMax and CMod outputs (fig. 8) are correlated with shift instability of RMax (fig. 9, top). This is in line with (112) and (114) in Theorems 2 and 3. Note that CMod outputs are nearly shift invariant regardless the characteristic frequency $\theta_l^{(J)}$ (fig. 9, bottom), as predicted by Theorem 1 (39).

7 Conclusion

In this paper, we investigated the shift invariance properties captured by the max pooling operator, when applied to a convolution layer with Gabor-like kernels. We established a validity domain for near-shift invariance and confirmed our predictions through an experimental setting based on the dual-tree complex wavelet packet transform. Our results indicate that the CMod operator can serve as a stable proxy for RMax, extracting comparable features—except at certain filter frequencies, where potential degeneracies may arise after max pooling. In this context, a follow-up study applying these principles to real-world architectures was published as a conference paper, and is included as supplementary material for anonymization purposes (Authors, 2024). We emphasize that this companion paper is not a prior version of the present submission but rather an application of its principles.

A link was missing between real- and complex-valued convolutions in CNNs. By comparing the outputs of CMod and RMax operators, we established a connection between these two worlds, creating opportunities for extensions of the results obtained for complex wavelet transforms. To paraphrase Tygert et al. (2016),

the correspondence between standard real-valued CNNs (using max pooling) and complex wavelets is no longer “just a vague analogy.”

References

- Krishna B. Athreya and Soumendra N. Lahiri. *Measure Theory and Probability Theory*, volume 19. Springer, 2006.
- Authors. Follow-up work, published as a conference paper. In *Anonymized to comply with the double-blind review process*, 2024.
- Aharon Azulay and Yair Weiss. Why do deep convolutional networks generalize so poorly to small image transformations? *Journal of Machine Learning Research*, 20(184):1–25, 2019.
- Radu Balan, Maneesh Singh, and Dongmian Zou. Lipschitz Properties for Deep Convolutional Networks. *Contemporary Mathematics*, 706:129–151, 2018.
- Ilker Bayram and Ivan W. Selesnick. On the Dual-Tree Complex Wavelet Packet and M-Band Transforms. *IEEE Transactions on Signal Processing*, 56(6):2298–2310, 2008.
- Alberto Bietti. Approximation and Learning with Deep Convolutional Models: A Kernel Perspective. In *International Conference on Learning Representations*, 2022.
- Alberto Bietti and Julien Mairal. Invariance and Stability of Deep Convolutional Representations. In *Advances in Neural Information Processing Systems*, pp. 1622, 2017.
- Alberto Bietti and Julien Mairal. Group invariance, stability to deformations, and complexity of deep convolutional representations. *The Journal of Machine Learning Research*, 20(1):876–924, 2019a.
- Alberto Bietti and Julien Mairal. On the Inductive Bias of Neural Tangent Kernels. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019b.
- Joan Bruna and Stéphane Mallat. Invariant Scattering Convolution Networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1872–1886, 2013.
- Jameson Cahill, Joseph W. Iverson, Dustin G. Mixon, and Daniel Packer. Group-Invariant Max Filtering. *Foundations of Computational Mathematics*, 2024.
- Zhengdao Chen. Multi-Layer Neural Networks as Trainable Ladders of Hilbert Spaces. In *Proceedings of the 40th International Conference on Machine Learning*, pp. 4294–4329, 2023.
- Fergal Cotter and Nick Kingsbury. A Learnable Scatnet: Locally Invariant Convolutional Layers. In *2019 IEEE International Conference on Image Processing (ICIP)*, pp. 350–354, 2019.
- Wojciech Czaja and Weilin Li. Analysis of time-frequency scattering transforms. *Applied and Computational Harmonic Analysis*, 47(1):149–171, 2019.
- Wojciech Czaja and Weilin Li. Rotationally Invariant Time–Frequency Scattering Transforms. *Journal of Fourier Analysis and Applications*, 26(1):4, 2020.
- Wojciech Czaja, Brandon Kolstoe, and David Korolov. Translation-Invariant Behavior of General Scattering Transforms. *arXiv:2410.21030*, 2024.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*. arXiv, 2021.
- Fernando Gama, Alejandro Ribeiro, and Joan Bruna. Stability of Graph Scattering Transforms. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.

- Shanel Gauthier, Benjamin Thérien, Laurent Alsène-Racicot, Muawiz Chaudhary, Irina Rish, Eugene Belilovsky, Michael Eickenberg, and Guy Wolf. Parametric Scattering Networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5749–5758, 2022.
- Kavya Gupta, Fateh Kaakai, Beatrice Pesquet-Popescu, Jean-Christophe Pesquet, and Fragkiskos D. Malliaros. Multivariate Lipschitz Analysis of the Stability of Neural Networks. *Frontiers in Signal Processing*, 2, 2022.
- Paul R. Halmos. *Measure Theory*. Springer, 2013. ISBN 978-1-4684-9440-2.
- Ali Hassani, Steven Walton, Nikhil Shah, Abulikemu Abuduweili, Jiachen Li, and Humphrey Shi. Escaping the Big Data Paradigm with Compact Transformers. *arXiv:2104.05704*, 2022.
- J.P. Havlicek, J.W. Havlicek, and A.C. Bovik. The analytic image. In *Proceedings of International Conference on Image Processing*, volume 2, pp. 446–449 vol.2, 1997.
- Ke Huang and Selin Aviyente. Wavelet feature selection for image classification. *IEEE Transactions on Image Processing*, 17(9):1709–1720, 2008.
- Nick Kingsbury. Complex wavelets for shift invariant analysis and filtering of signals. *Applied and computational harmonic analysis*, 10(3):234–253, 2001.
- Nick Kingsbury. Design of Q-shift complex wavelets for image processing using frequency domain energy minimization. In *Proceedings International Conference on Image Processing*, volume 1, pp. I–1013, 2003.
- Nick Kingsbury and Julian Magarey. Wavelet Transforms in Image Processing. In *Signal Analysis and Prediction*, Applied and Numerical Harmonic Analysis, pp. 27–46. Birkhäuser, Boston, MA, 1998. ISBN 978-1-4612-1768-8.
- Andrew Laine and Jian Fan. Texture Classification by Wavelet Packet Signatures. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(11):1186–1191, 1993.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2323, 1998.
- Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.
- Kunchang Li, Yali Wang, Junhao Zhang, Peng Gao, Guanglu Song, Yu Liu, Hongsheng Li, and Yu Qiao. UniFormer: Unifying Convolution and Self-Attention for Visual Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(10):12581–12600, 2023.
- Bin Liao and Fen Peng. Rotation-invariant texture features extraction using Dual-Tree Complex Wavelet Transform. In *2010 International Conference on Information, Networking and Automation (ICINA)*, volume 1, 2010.
- Stéphane Mallat. *A Wavelet Tour of Signal Processing : The Sparse Way*. Academic Press, 2009.
- Stéphane Mallat. Group invariant scattering. *Communications on Pure and Applied Mathematics*, 65(10):1331–1398, 2012.
- Stéphane Mallat. Understanding deep convolutional networks. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 374(2065):20150203, 2016.
- Yves Meyer. Principe d’incertitude, bases hilbertiennes et algèbres d’opérateurs. In *Séminaire Bourbaki*, volume 662. Société Mathématique de France, 1985.
- G. P. Nason and B. W. Silverman. The Stationary Wavelet Transform and some Statistical Applications. In *Wavelets and Statistics*, Lecture Notes in Statistics, pp. 281–299. Springer, 1995.
- Edouard Oyallon. *Analyzing and Introducing Structures in Deep Convolutional Neural Networks*. Doctoral thesis, Paris Sciences et Lettres, 2017.

- Edouard Oyallon, Eugene Belilovsky, Sergey Zagoruyko, and Michal Valko. Compressing the Input for CNNs with the First-Order Scattering Transform. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- Kun Il Park and M. Park. *Fundamentals of Probability and Stochastic Processes with Applications to Communications*. Springer, 2018.
- Juan C. Pérez, Motasem Alfarrar, Guillaume Jeanneret, Adel Bibi, Ali Thabet, Bernard Ghanem, and Pablo Arbeláez. Gabor Layers Enhance Network Robustness. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. ISBN 978-3-030-58545-7.
- Michael Perlmutter, Feng Gao, Guy Wolf, and Matthew Hirn. Geometric Wavelet Scattering Networks on Compact Riemannian Manifolds. In *Proceedings of The First Mathematical and Scientific Machine Learning Conference*, pp. 570–604. PMLR, 2020.
- Stefan Pittner and Sagar V. Kamarthi. Feature extraction from wavelet coefficients for pattern recognition tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(1):83–88, 1999.
- Mehang Rai and Pablo Rivas. A Review of Convolutional Neural Networks and Gabor Filters in Object Recognition. In *2020 International Conference on Computational Science and Computational Intelligence (CSCI)*, pp. 1560–1567, 2020.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.
- Meyer Scetbon and Zaid Harchaoui. Harmonic Decompositions of Convolutional Networks. In *International Conference on Machine Learning*, pp. 8522–8532. PMLR, 2020.
- Ivan W. Selesnick, Richard Baraniuk, and Nick Kingsbury. The dual-tree complex wavelet transform. *IEEE Signal Processing Magazine*, 22(6):123–151, 2005.
- C.E. Shannon. Communication in the Presence of Noise. *Proceedings of the IRE*, 37(1):10–21, 1949.
- Laurent Sifre and Stéphane Mallat. Rotation, Scaling and Deformation Invariant Scattering for Texture Discrimination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1233–1240, 2013.
- Amarjot Singh and Nick Kingsbury. Dual-Tree wavelet scattering network with parametric log transformation for object classification. In *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, 2017.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv:1312.6199 [cs]*, 2014.
- Antonio Torralba and Aude Oliva. Statistics of natural image categories. *Network: Computation in Neural Systems*, 14(3):391–412, 2003.
- Mark Tygert, Joan Bruna, Soumith Chintala, Yann LeCun, Serkan Piantino, and Arthur Szlam. A Mathematical Motivation for Complex-Valued Convolutional Networks. *Neural Computation*, 28(5):815–825, 2016.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, pp. 5998–6008, 2017.
- Martin Vetterli. Wavelets, approximation, and compression. *IEEE Signal Processing Magazine*, 18(5):59–73, 2001.

- Aladin Virmaux and Kevin Scaman. Lipschitz regularity of deep neural networks: Analysis and efficient estimation. In *Advances in Neural Information Processing Systems*, 2018.
- Irène Waldspurger. *Wavelet Transform Modulus : Phase Retrieval and Scattering*. Doctoral thesis, Ecole normale supérieure, Paris, 2015.
- Thomas Wiatowski and Helmut Bölcskei. A Mathematical Theory of Deep Convolutional Neural Networks for Feature Extraction. *IEEE Transactions on Information Theory*, 64(3):1845–1866, 2018.
- Haiping Wu, Bin Xiao, Noel Codella, Mengchen Liu, Xiyang Dai, Lu Yuan, and Lei Zhang. CvT: Introducing Convolutions to Vision Transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 22–31, 2021.
- Gary G. Yen. Wavelet packet feature extraction for vibration monitoring. *IEEE Transactions on Industrial Electronics*, 47(3):650–667, 2000.
- Lirong Yin, Lei Wang, Siyu Lu, Ruiyang Wang, Youshuai Yang, Bo Yang, Shan Liu, Ahmed AlSanad, Salman A. AlQahtani, Zhengtong Yin, Xiaolu Li, Xiaobing Chen, and Wenfeng Zheng. Convolution-Transformer for Image Feature Extraction. *Computer Modeling in Engineering & Sciences*, 141(1):87–106, 2024.
- Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep neural networks? In *Advances in Neural Information Processing Systems*, pp. 3320–3328, 2014.
- Kun Yuan, Shaopeng Guo, Ziwei Liu, Aojun Zhou, Fengwei Yu, and Wei Wu. Incorporating Convolution Designs Into Visual Transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 579–588, 2021.
- John Zarka, Louis Thiry, Tomás Angles, and Stéphane Mallat. Deep Network Classification by Scattering and Homotopy Dictionary Learning. In *International Conference on Learning Representations*, 2020.
- John Zarka, Florentin Guth, and Stéphane Mallat. Separation and Concentration in Deep Networks. In *International Conference on Learning Representations*, 2021.
- Richard Zhang. Making Convolutional Networks Shift-Invariant Again. In *International Conference on Machine Learning*, 2019.
- Dongmian Zou and Gilad Lerman. Graph convolutional neural networks via scattering. *Applied and Computational Harmonic Analysis*, 49(3):1046–1074, 2020.
- Dongmian Zou, Radu Balan, and Maneesh Singh. On Lipschitz Bounds of General Convolutional Neural Networks. *IEEE Transactions on Information Theory*, 66(3):1738–1759, 2020.
- Xueyan Zou, Fanyi Xiao, Zhiding Yu, Yuheng Li, and Yong Jae Lee. Delving Deeper into Anti-Aliasing in ConvNets. *International Journal of Computer Vision*, 131(1):67–81, 2023.
- Monty-Maximilian Zühlke and Daniel Kudenko. Adversarial Robustness of Neural Networks From the Perspective of Lipschitz Calculus: A Survey. *ACM Comput. Surv.*, 2024.

A Appendix – Proofs

A.1 Proof of Lemma 1

Proof. Applying the Fourier transform on (15) yields, for any $\xi \in \mathbb{R}^2$,

$$\widehat{F}_0(\xi) = (\widehat{F * \Psi})(\xi - \nu) = \mathcal{T}_\nu(\widehat{F}\widehat{\Psi})(\xi). \quad (141)$$

By hypothesis on Ψ , we have

$$\text{supp}(\widehat{F}\widehat{\Psi}) \subset \text{supp}\widehat{\Psi} \subset B_\infty(-\nu, \varepsilon/2). \quad (142)$$

The translation operator \mathcal{T}_ν shifts the support with respect to ν , which yields (16). \square

A.2 Proof of Proposition 1

Proof. Using the 2D Plancherel formula, we compute

$$\begin{aligned}
\|\mathcal{T}_{\mathbf{h}}F_0 - F_0\|_{L^2}^2 &= \frac{1}{4\pi^2} \left\| \widehat{\mathcal{T}_{\mathbf{h}}F_0} - \widehat{F_0} \right\|_{L^2}^2 \\
&= \frac{1}{4\pi^2} \iint_{B_\infty(\varepsilon/2)} \left| \widehat{F_0}(\boldsymbol{\xi}) \right|^2 \left| e^{-i\langle \mathbf{h}, \boldsymbol{\xi} \rangle} - 1 \right|^2 d^2\boldsymbol{\xi} \\
&= \frac{1}{4\pi^2} \iint_{B_\infty(\varepsilon/2)} \left| \widehat{F_0}(\boldsymbol{\xi}) \right|^2 (2 - 2\cos \langle \mathbf{h}, \boldsymbol{\xi} \rangle) d^2\boldsymbol{\xi} \\
&\leq \frac{1}{4\pi^2} \iint_{B_\infty(\varepsilon/2)} \left| \widehat{F_0}(\boldsymbol{\xi}) \right|^2 |\langle \mathbf{h}, \boldsymbol{\xi} \rangle|^2 d^2\boldsymbol{\xi},
\end{aligned}$$

because $\cos t \geq 1 - \frac{t^2}{2}$. Note that the integral is computed on a compact domain because, according to Lemma 1, $\text{supp } \widehat{F_0} \subset B_\infty(\varepsilon/2)$. Next, we use the Cauchy-Schwarz inequality to compute:

$$\begin{aligned}
\forall \boldsymbol{\xi} \in B_\infty(\varepsilon/2), \quad |\langle \mathbf{h}, \boldsymbol{\xi} \rangle| &\leq \|\mathbf{h}\|_1 \cdot \|\boldsymbol{\xi}\|_\infty \\
&\leq \frac{\varepsilon}{2} \|\mathbf{h}\|_1.
\end{aligned}$$

Therefore,

$$\|\mathcal{T}_{\mathbf{h}}F_0 - F_0\|_{L^2}^2 \leq \frac{\varepsilon}{4} \|\mathbf{h}\|_1^2 \|F_0\|_{L^2}^2, \quad (143)$$

which yields the result. \square

A.3 Proof of Lemma 2

Proof. Since $F \in \mathcal{V}^{(s)}$, the two-dimensional version of Shannon's sampling theorem (Mallat, 2009, Theorem 3.11, p. 81) yields

$$F = \sum_{\mathbf{n} \in \mathbb{Z}^2} X[\mathbf{n}] \phi_{\mathbf{n}}^{(s)}, \quad \text{and} \quad \widehat{F} = \sum_{\mathbf{n} \in \mathbb{Z}^2} X[\mathbf{n}] \widehat{\phi}_{\mathbf{n}}^{(s)}. \quad (144)$$

Moreover, using (19), we can show that, for any $\boldsymbol{\xi} \in B_\infty(\pi/s)$,

$$\widehat{\phi}_{\mathbf{n}}^{(s)}(\boldsymbol{\xi}) = \widehat{\phi}^{(s)}(\boldsymbol{\xi}) e^{-i\langle s\boldsymbol{\xi}, \mathbf{n} \rangle} = s e^{-i\langle s\boldsymbol{\xi}, \mathbf{n} \rangle}. \quad (145)$$

Therefore, plugging (145) into (144) proves (22).

Then, by combining (22) with the Plancherel formula, we get

$$\begin{aligned}
\|F\|_{L^2}^2 &= \frac{1}{4\pi^2} \|\widehat{F}\|_{L^2}^2 \\
&= \frac{1}{4\pi^2} \iint_{B_\infty(\pi/s)} |\widehat{F}(\boldsymbol{\xi})|^2 d^2\boldsymbol{\xi} \\
&= \frac{1}{4\pi^2} \iint_{B_\infty(\pi/s)} |s \widehat{X}(s\boldsymbol{\xi})|^2 d^2\boldsymbol{\xi}.
\end{aligned}$$

The integral is performed on $B_\infty(\pi/s)$ because $F \in \mathcal{V}^{(s)}$. Then, by applying the change of variable $\boldsymbol{\xi}' \leftarrow s\boldsymbol{\xi}$, we get

$$\begin{aligned}
\|F\|_{L^2}^2 &= \frac{1}{4\pi^2} \iint_{B_\infty(\pi)} |\widehat{X}(\boldsymbol{\xi}')|^2 d^2\boldsymbol{\xi}' \\
&= \frac{1}{4\pi^2} \|\widehat{X}\|_{L^2}^2 = \|X\|_2^2,
\end{aligned}$$

hence (23), which concludes the proof. \square

A.4 Proof of Proposition 2

Proof. First, F_X and Ψ_W are well defined because $X \in l_{\mathbb{R}}^2(\mathbb{Z}^2)$ and $W \in l_{\mathbb{C}}^2(\mathbb{Z}^2)$. By construction, F_X and $\Psi_W \in \mathcal{V}^{(s)}$. Therefore, according to Shannon's sampling theorem (Mallat, 2009, Theorem 3.11, p. 81),

$$F_X := s \sum_{\mathbf{n} \in \mathbb{Z}^2} F_X(s\mathbf{n}) \Phi_{\mathbf{n}}^{(s)} \quad \text{and} \quad \Psi_W := s \sum_{\mathbf{n} \in \mathbb{Z}^2} \Psi_W(s\mathbf{n}) \Phi_{\mathbf{n}}^{(s)}. \quad (146)$$

By uniqueness of decompositions in an orthonormal basis, we get (26). Moreover, using (22) in Lemma 2, we get, for any $\boldsymbol{\xi} \in B_{\infty}(\pi/s)$,

$$\widehat{\Psi}_W(\boldsymbol{\xi}) = s \widehat{W}(s\boldsymbol{\xi}). \quad (147)$$

Since $\widehat{\Psi}_W(\boldsymbol{\xi}) = 0$ outside $B_{\infty}(\pi/s)$, (147) is true for any $\boldsymbol{\xi} \in \mathbb{R}^2$. Therefore, by hypothesis on W ,

$$\text{supp } \widehat{\Psi}_W \subset B_{\infty}(\boldsymbol{\theta}/s, \kappa/(2s)), \quad (148)$$

which yields (25).

We now prove (27). For $\mathbf{n} \in \mathbb{Z}^2$, we compute:

$$\begin{aligned} (F_X * \bar{\Psi}_W)(ms\mathbf{n}) &= \iint_{\mathbb{R}^2} F_X(ms\mathbf{n} - \mathbf{x}) \bar{\Psi}_W(\mathbf{x}) d^2\mathbf{x} \\ &= \iint_{\mathbb{R}^2} \sum_{\mathbf{p} \in \mathbb{Z}^2} X[\mathbf{p}] \Phi_{\mathbf{p}}^{(s)}(ms\mathbf{n} - \mathbf{x}) \bar{\Psi}_W(\mathbf{x}) d^2\mathbf{x} \\ &= \sum_{\mathbf{p} \in \mathbb{Z}^2} X[\mathbf{p}] \iint_{\mathbb{R}^2} \Phi_{\mathbf{p}}^{(s)}(ms\mathbf{n} - \mathbf{x}) \bar{\Psi}_W(\mathbf{x}) d^2\mathbf{x}. \end{aligned}$$

The sum-integral interchange is possible because X has a finite support. Then:

$$(F_X * \bar{\Psi}_W)(ms\mathbf{n}) = \sum_{\mathbf{p} \in \mathbb{Z}^2} X[\mathbf{p}] \iint_{\mathbb{R}^2} \bar{\Psi}_W(\mathbf{x}) \Phi^{(s)}(s(m\mathbf{n} - \mathbf{p}) - \mathbf{x}) d^2\mathbf{x} \quad (149)$$

$$= \sum_{\mathbf{p} \in \mathbb{Z}^2} X[\mathbf{p}] (\bar{\Psi}_W * \Phi^{(s)})(s(m\mathbf{n} - \mathbf{p})) \quad (150)$$

Since $\{\Phi_{\mathbf{n}}^{(s)}\}_{\mathbf{n} \in \mathbb{Z}^2}$ is an orthonormal basis of $\mathcal{V}^{(s)}$, the definition of Ψ_W in (24) implies, for any $\mathbf{p}' \in \mathbb{Z}^2$,

$$\bar{W}[\mathbf{p}'] = \langle \Psi_W, \Phi_{-\mathbf{p}'}^{(s)} \rangle = (\bar{\Psi}_W * \Phi^{(s)})(s\mathbf{p}'), \quad (151)$$

because $\Phi^{(s)}$ is even. Therefore, plugging (151) with $\mathbf{p}' \leftarrow (m\mathbf{n} - \mathbf{p})$ into (150) yields

$$(F_X * \bar{\Psi}_W)(ms\mathbf{n}) = \sum_{\mathbf{p} \in \mathbb{Z}^2} X[\mathbf{p}] \bar{W}[m\mathbf{n} - \mathbf{p}] = (X * \bar{W})[m\mathbf{n}], \quad (152)$$

hence the result. \square

A.5 Proof of Lemma 3

Proof. Let $\mathbf{u} \in \mathbb{R}^2$. By definition of $F_{\mathcal{T}_{\mathbf{u}}X}$ and $\mathcal{T}_{\mathbf{u}}X$,

$$F_{\mathcal{T}_{\mathbf{u}}X} = s \sum_{\mathbf{n} \in \mathbb{Z}^2} \mathcal{T}_{s\mathbf{u}} F_X(s\mathbf{n}) \Phi_{\mathbf{n}}^{(s)}. \quad (153)$$

On the other hand, $F_X \in \mathcal{V}^{(s)}$ by construction. Therefore, $\mathcal{T}_{s\mathbf{u}} F_X \in \mathcal{V}^{(s)}$. Then, according to Shannon's sampling theorem (Mallat, 2009, Theorem 3.11, p. 81), we get

$$\mathcal{T}_{s\mathbf{u}} F_X = s \sum_{\mathbf{n} \in \mathbb{Z}^2} \mathcal{T}_{s\mathbf{u}} F_X(s\mathbf{n}) \Phi_{\mathbf{n}}^{(s)}, \quad (154)$$

which concludes the proof. \square

A.6 Proof of Lemma 4

Proof. Let us write:

$$\sum_{\mathbf{n} \in \mathbb{Z}^2} |\mathcal{T}_{\mathbf{h}} F_0(s' \mathbf{n}) - F_0(s' \mathbf{n})|^2 = \sum_{\mathbf{n} \in \mathbb{Z}^2} |F_1(s' \mathbf{n})|^2 = \frac{1}{s'^2} \|X_1\|_2^2, \quad (155)$$

where we have denoted, for any $\mathbf{n} \in \mathbb{Z}^2$,

$$F_1 := \mathcal{T}_{\mathbf{h}} F_0 - F_0 \quad \text{and} \quad X_1[\mathbf{n}] := s' F_1(s' \mathbf{n}). \quad (156)$$

According to Proposition 2 (25), $\Psi_{\mathbf{W}} \in \mathcal{V}(\boldsymbol{\theta}/s, \kappa/s)$. Therefore, according to Lemma 1,

$$\text{supp } \widehat{F_0} \subset B_{\infty} \left(\frac{\kappa}{2s} \right). \quad (157)$$

Moreover, by hypothesis, $\kappa \leq \pi/m$; thus,

$$B_{\infty} \left(\frac{\kappa}{2s} \right) \subset B_{\infty} \left(\frac{\pi}{s'} \right), \quad (158)$$

which yields (35), and $F_1 \in \mathcal{V}(s')$. Then, according to Lemma 2 (23) with $X \leftarrow X_1$, $F \leftarrow F_1$ and $s \leftarrow s'$,

$$\|X_1\|_2 = \|F_1\|_{L^2} = \|\mathcal{T}_{\mathbf{h}} F_0 - F_0\|_{L^2}. \quad (159)$$

Therefore, plugging (159) into (155) yields (36).

Furthermore, according again to Lemma 2,

$$\|F_0\|_{L^2}^2 = \|X_0\|_2^2, \quad (160)$$

where we have defined, for any $\mathbf{n} \in \mathbb{Z}^2$,

$$X_0[\mathbf{n}] := s' F_0(s' \mathbf{n}). \quad (161)$$

Then,

$$\begin{aligned} \|X_0\|_2^2 &= s'^2 \sum_{\mathbf{n} \in \mathbb{Z}^2} |(F_X * \bar{\Psi}_{\mathbf{W}})(s' \mathbf{n})|^2 && \text{(acc. to (34))} \\ &= s'^2 \sum_{\mathbf{n} \in \mathbb{Z}^2} |(X * \bar{\mathbf{W}}) \downarrow (2m)[\mathbf{n}]|^2 && \text{(acc. to Proposition 2 with } m \leftarrow 2m) \\ &= s'^2 \|U_m^{\text{mod}} X\|_2^2. && \text{(acc. to (5))} \end{aligned}$$

Finally, plugging this result into (160) yields (37) and concludes the proof. \square

A.7 Proof of Theorem 1

Proof. As in Lemma 4, we consider the low-frequency function F_0 satisfying (34), and denote $s' := 2ms$. We can write

$$|F_X * \bar{\Psi}_{\mathbf{W}}| = |F_0| \quad \text{and} \quad |\mathcal{T}_{s\mathbf{u}} F_X * \bar{\Psi}_{\mathbf{W}}| = |\mathcal{T}_{s\mathbf{u}} F_0|. \quad (162)$$

Recall that $U_m^{\text{mod}} X = |(X * \bar{\mathbf{W}}) \downarrow (2m)|$, such as defined in (5). According to Proposition 2 (27) and Corollary 1 (32) with $m \leftarrow 2m$, we therefore get

$$U_m^{\text{mod}} X[\mathbf{n}] = |F_0(s' \mathbf{n})|; \quad (163)$$

$$U_m^{\text{mod}} (\mathcal{T}_{\mathbf{u}} X)[\mathbf{n}] = |(\mathcal{T}_{s\mathbf{u}} F_0)(s' \mathbf{n})|. \quad (164)$$

Then, using (163), (164) and the reverse triangle inequality,

$$\begin{aligned} \|U_m^{\text{mod}}(\mathcal{T}_{\mathbf{u}}\mathbf{X}) - U_m^{\text{mod}}\mathbf{X}\|_2^2 &= \sum_{\mathbf{n} \in \mathbb{Z}^2} \left| |(\mathcal{T}_{s\mathbf{u}}F_0)(s'\mathbf{n})| - |F_0(s'\mathbf{n})| \right|^2 \\ &\leq \sum_{\mathbf{n} \in \mathbb{Z}^2} \left| (\mathcal{T}_{s\mathbf{u}}F_0)(s'\mathbf{n}) - F_0(s'\mathbf{n}) \right|^2. \end{aligned}$$

Since condition (38) is satisfied, we can use Lemma 4 (36) with $\mathbf{h} \leftarrow s\mathbf{u}$:

$$\|U_m^{\text{mod}}(\mathcal{T}_{\mathbf{u}}\mathbf{X}) - U_m^{\text{mod}}\mathbf{X}\|_2^2 \leq \frac{1}{s'^2} \|\mathcal{T}_{s\mathbf{u}}F_0 - F_0\|_{L^2}^2 \quad (165)$$

Next, according to Proposition 1 with $\varepsilon \leftarrow \kappa/s$ and $\mathbf{h} \leftarrow s\mathbf{u}$, we then get the following bound:

$$\|U_m^{\text{mod}}(\mathcal{T}_{\mathbf{u}}\mathbf{X}) - U_m^{\text{mod}}\mathbf{X}\|_2^2 \leq \frac{\alpha(\kappa\mathbf{u})^2}{s'^2} \|F_0\|_{L^2}^2. \quad (166)$$

Finally, using Lemma 4 (37) yields (39), which completes the proof. \square

A.8 Proof of Proposition 3

Proof. Let $\mathbf{X} \in l_{\mathbb{R}}^2(\mathbb{Z}^2)$ and $s > 0$. We consider $F_0 \in L_{\mathbb{C}}^2(\mathbb{R}^2)$ as the “low-frequency” function satisfying (34). Again, we introduce $s' := 2ms$ and $\mathbf{X}_0 \in l_{\mathbb{C}}^2(\mathbb{Z}^2)$ satisfying (161). Moreover, for any $\mathbf{Y} \in l_{\mathbb{R}}^2(\mathbb{Z}^2)$, we denote by $F_{\mathbf{Y}}^{(s')}$ the Shannon interpolation of \mathbf{Y} parameterized by s' , analogously to (24):

$$F_{\mathbf{Y}}^{(s')} := \sum_{\mathbf{n} \in \mathbb{Z}^2} \mathbf{Y}[\mathbf{n}] \Phi_{\mathbf{n}}^{(s')}. \quad (167)$$

On the one hand, Lemma 4 provides (37). On the other hand, we seek a similar result with $\mathbf{X} \leftarrow \mathcal{T}_{\mathbf{u}}\mathbf{X}$. For this purpose, (164) can be rewritten

$$U_m^{\text{mod}}(\mathcal{T}_{\mathbf{u}}\mathbf{X})[\mathbf{n}] = |\mathcal{T}_{s'\mathbf{u}'}F_0(s'\mathbf{n})|, \quad (168)$$

with $\mathbf{u}' := \mathbf{u}/(2m)$. Furthermore, according to Lemma 4 (35), $F_0 \in \mathcal{V}^{(s')}$. Therefore, Shannon’s sampling theorem (Mallat, 2009, Theorem 3.11, p. 81) with $s \leftarrow s'$ yields

$$\begin{aligned} F_0 &= s' \sum_{\mathbf{n} \in \mathbb{Z}^2} F_0(s'\mathbf{n}) \Phi_{\mathbf{n}}^{(s')} \\ &= \sum_{\mathbf{n} \in \mathbb{Z}^2} \mathbf{X}_0[\mathbf{n}] \Phi_{\mathbf{n}}^{(s')} = F_{\mathbf{X}_0}^{(s')}, \end{aligned}$$

where we have used the notations introduced in (161) and (167). Then, using Lemma 3 with $\mathbf{X} \leftarrow \mathbf{X}_0$, $\mathbf{u} \leftarrow \mathbf{u}'$ and $s \leftarrow s'$, we get

$$F_{\mathcal{T}_{\mathbf{u}'}\mathbf{X}_0}^{(s')} = \mathcal{T}_{s'\mathbf{u}'}F_{\mathbf{X}_0}^{(s')} = \mathcal{T}_{s'\mathbf{u}'}F_0. \quad (169)$$

Moreover, (26) (from Proposition 2) with $\mathbf{X} \leftarrow \mathcal{T}_{\mathbf{u}'}\mathbf{X}_0$ and $s \leftarrow s'$ becomes

$$\mathcal{T}_{\mathbf{u}'}\mathbf{X}_0[\mathbf{n}] = s' F_{\mathcal{T}_{\mathbf{u}'}\mathbf{X}_0}^{(s')}(s'\mathbf{n}), \quad (170)$$

and inserting (169) into (170) yields

$$\mathcal{T}_{\mathbf{u}'}\mathbf{X}_0[\mathbf{n}] = s' \mathcal{T}_{s'\mathbf{u}'}F_0(s'\mathbf{n}). \quad (171)$$

Therefore, (168) and (171) imply

$$\|U_m^{\text{mod}}(\mathcal{T}_{\mathbf{u}}\mathbf{X})\|_2 = \frac{1}{s'} \|\mathcal{T}_{\mathbf{u}'}\mathbf{X}_0\|_2. \quad (172)$$

Moreover, since $F_0 \in \mathcal{V}^{(s')}$, and according to (171), we can use Lemma 2 with $s \leftarrow s'$, $F \leftarrow \mathcal{T}_{s'\mathbf{u}'}F_0$ and $\mathbf{X} \leftarrow \mathcal{T}_{\mathbf{u}'}\mathbf{X}_0$. We get

$$\|\mathcal{T}_{\mathbf{u}'}\mathbf{X}_0\|_2 = \|\mathcal{T}_{s'\mathbf{u}'}F_0\|_{L^2} = \|F_0\|_{L^2}, \quad (173)$$

and plugging (173) into (172) yields

$$\|U_m^{\text{mod}}(\mathcal{T}_{\mathbf{u}'}\mathbf{X})\|_2 = \frac{1}{s'} \|F_0\|_{L^2}. \quad (174)$$

Finally, considering Lemma 4 (37) concludes the proof. \square

A.9 Proof of Proposition 4

Proof. Let us write:

$$\begin{aligned} & (F * \text{Re}\bar{\Psi})(\mathbf{x} + \mathbf{h}) - |(F * \bar{\Psi})(\mathbf{x})| G(\mathbf{x}, \mathbf{h}) \\ &= \text{Re}((F * \bar{\Psi})(\mathbf{x} + \mathbf{h})) - |(F * \bar{\Psi})(\mathbf{x})| \text{Re}\left(e^{-i\langle \boldsymbol{\nu}, \mathbf{h} \rangle} e^{iH(\mathbf{x})}\right) \\ &= \text{Re}((F * \bar{\Psi})(\mathbf{x} + \mathbf{h})) - \text{Re}\left(|(F * \bar{\Psi})(\mathbf{x})| e^{iH(\mathbf{x})} e^{-i\langle \boldsymbol{\nu}, \mathbf{h} \rangle}\right) \\ &= \text{Re}((F * \bar{\Psi})(\mathbf{x} + \mathbf{h})) - \text{Re}\left((F * \bar{\Psi})(\mathbf{x}) e^{-i\langle \boldsymbol{\nu}, \mathbf{h} \rangle}\right) \\ &= \text{Re}\left((F * \bar{\Psi})(\mathbf{x} + \mathbf{h}) - (F * \bar{\Psi})(\mathbf{x}) e^{-i\langle \boldsymbol{\nu}, \mathbf{h} \rangle}\right). \end{aligned}$$

Therefore,

$$\begin{aligned} \left| (F * \text{Re}\bar{\Psi})(\mathbf{x} + \mathbf{h}) - |(F * \bar{\Psi})(\mathbf{x})| G(\mathbf{x}, \mathbf{h}) \right| &\leq \left| (F * \bar{\Psi})(\mathbf{x} + \mathbf{h}) - (F * \bar{\Psi})(\mathbf{x}) e^{-i\langle \boldsymbol{\nu}, \mathbf{h} \rangle} \right| \\ &= \left| F_0(\mathbf{x} + \mathbf{h}) e^{-i\langle \boldsymbol{\nu}, \mathbf{x} + \mathbf{h} \rangle} - F_0(\mathbf{x}) e^{-i\langle \boldsymbol{\nu}, \mathbf{x} + \mathbf{h} \rangle} \right|, \end{aligned}$$

which yields (44) and concludes the proof. \square

A.10 Proof of Lemma 5

Proof. We apply Proposition 4 with $\mathbf{h} \leftarrow \mathbf{h}_n^{\max}$ and $\mathbf{h} \leftarrow \mathbf{h}'_n^{\max}$, respectively:

$$A_X^{\max}(\mathbf{x}_n) \leq \tilde{A}_X(\mathbf{x}_n, \mathbf{h}_n^{\max}) + |F_0(\mathbf{x}_n + \mathbf{h}_n^{\max}) - F_0(\mathbf{x}_n)|; \quad (175)$$

$$\tilde{A}_X^{\max}(\mathbf{x}_n) \leq A_X(\mathbf{x}_n, \mathbf{h}'_n^{\max}) + |F_0(\mathbf{x}_n + \mathbf{h}'_n^{\max}) - F_0(\mathbf{x}_n)|. \quad (176)$$

By construction, we have, for any $\mathbf{p} \in \{-q, \dots, q\}^2$,

$$\tilde{A}_X(\mathbf{x}_n, \mathbf{h}_{\mathbf{p}}) \leq \tilde{A}_X^{\max}(\mathbf{x}_n) \quad \text{and} \quad A_X(\mathbf{x}_n, \mathbf{h}_{\mathbf{p}}) \leq A_X^{\max}(\mathbf{x}_n). \quad (177)$$

Moreover, by definition, there exists \mathbf{p} and $\mathbf{p}' \in \{-q, \dots, q\}^2$ such that $\mathbf{h}_n^{\max} = \mathbf{h}_{\mathbf{p}}$ and $\mathbf{h}'_n^{\max} = \mathbf{h}_{\mathbf{p}'}$. Therefore, (175) and (176) yield, respectively,

$$A_X^{\max}(\mathbf{x}_n) \leq \tilde{A}_X^{\max}(\mathbf{x}_n) + |F_0(\mathbf{x}_n + \mathbf{h}_n^{\max}) - F_0(\mathbf{x}_n)|; \quad (178)$$

$$\tilde{A}_X^{\max}(\mathbf{x}_n) \leq A_X^{\max}(\mathbf{x}_n) + |F_0(\mathbf{x}_n + \mathbf{h}'_n^{\max}) - F_0(\mathbf{x}_n)|, \quad (179)$$

which yields (68) and concludes the proof. \square

A.11 Proof of Proposition 5

Proof. Let us write:

$$\begin{aligned}
\|U_m^{\text{mod}}\mathbf{X} - U_{m,q}^{\text{max}}\mathbf{X}\|_2^2 &= \sum_{\mathbf{n} \in \mathbb{Z}^2} (U_m^{\text{mod}}\mathbf{X}[\mathbf{n}] - U_{m,q}^{\text{max}}\mathbf{X}[\mathbf{n}])^2 \\
&= \sum_{\mathbf{n} \in \mathbb{Z}^2} \left(U_m^{\text{mod}}\mathbf{X}[\mathbf{n}] - U_m^{\text{mod}}\mathbf{X}[\mathbf{n}] \max_{\|\mathbf{p}\|_\infty \leq q} G_{\mathbf{X}}(\mathbf{x}_{\mathbf{n}}, \mathbf{h}_{\mathbf{p}}) \right. \\
&\quad \left. + U_m^{\text{mod}}\mathbf{X}[\mathbf{n}] \max_{\|\mathbf{p}\|_\infty \leq q} G_{\mathbf{X}}(\mathbf{x}_{\mathbf{n}}, \mathbf{h}_{\mathbf{p}}) - U_{m,q}^{\text{max}}\mathbf{X}[\mathbf{n}] \right)^2 \\
&= \sum_{\mathbf{n} \in \mathbb{Z}^2} \left(\delta_{m,q}\mathbf{X}[\mathbf{n}] + \tilde{A}_{\mathbf{X}}^{\text{max}}(\mathbf{x}_{\mathbf{n}}) - A_{\mathbf{X}}^{\text{max}}(\mathbf{x}_{\mathbf{n}}) \right)^2,
\end{aligned}$$

according to (60), (65) and (66). Then, using the triangle inequality, we get

$$\|U_m^{\text{mod}}\mathbf{X} - U_{m,q}^{\text{max}}\mathbf{X}\|_2 \leq \|\delta_{m,q}\mathbf{X}\|_2 + \left(\sum_{\mathbf{n} \in \mathbb{Z}^2} \left(\tilde{A}_{\mathbf{X}}^{\text{max}}(\mathbf{x}_{\mathbf{n}}) - A_{\mathbf{X}}^{\text{max}}(\mathbf{x}_{\mathbf{n}}) \right)^2 \right)^{1/2}. \quad (180)$$

Furthermore, Lemma 5 yields

$$\sum_{\mathbf{n} \in \mathbb{Z}^2} \left(\tilde{A}_{\mathbf{X}}^{\text{max}}(\mathbf{x}_{\mathbf{n}}) - A_{\mathbf{X}}^{\text{max}}(\mathbf{x}_{\mathbf{n}}) \right)^2 \leq \sum_{\mathbf{n} \in \mathbb{Z}^2} \max_{\mathbf{h} \in \{\mathbf{h}_{\mathbf{n}}^{\text{max}}, \mathbf{h}'_{\mathbf{n}}\}} \left| F_0(\mathbf{x}_{\mathbf{n}} + \mathbf{h}) - F_0(\mathbf{x}_{\mathbf{n}}) \right|^2 \quad (181)$$

$$\leq \sum_{\mathbf{n} \in \mathbb{Z}^2} \left| F_0(\mathbf{x}_{\mathbf{n}} + \mathbf{h}_0) - F_0(\mathbf{x}_{\mathbf{n}}) \right|^2, \quad (182)$$

according to Hypothesis 1. Next, since (38) is satisfied, we can use Lemma 4 (36) with $\mathbf{h} \leftarrow \mathbf{h}_0$. We get

$$\begin{aligned}
\sum_{\mathbf{n} \in \mathbb{Z}^2} \left(\tilde{A}_{\mathbf{X}}^{\text{max}}(\mathbf{x}_{\mathbf{n}}) - A_{\mathbf{X}}^{\text{max}}(\mathbf{x}_{\mathbf{n}}) \right)^2 &\leq \frac{1}{4m^2s^2} \|\mathcal{T}_{\mathbf{h}_0}F_0 - F_0\|_{L^2}^2 \\
&\leq \alpha(\kappa\mathbf{h}_0/s)^2 \frac{1}{4m^2s^2} \|F_0\|_{L^2}^2 \quad (\text{acc. to Proposition 1}) \\
&= \alpha(\kappa\mathbf{h}_0/s)^2 \|U_m^{\text{mod}}\mathbf{X}\|_2^2. \quad (\text{acc. to Lemma 4 (37)})
\end{aligned}$$

Since, according to Hypothesis 1, $\|\mathbf{h}_0\|_2 = \sqrt{2}qms$, it comes that $\|\mathbf{h}_0\|_1 = 2qms$. Therefore,

$$\alpha(\kappa\mathbf{h}_0/s)^2 = \frac{\kappa^2 \|\mathbf{h}_0\|_1^2}{4s^2} = (qm\kappa)^2, \quad (183)$$

which yields

$$\sum_{\mathbf{n} \in \mathbb{Z}^2} \left(\tilde{A}_{\mathbf{X}}^{\text{max}}(\mathbf{x}_{\mathbf{n}}) - A_{\mathbf{X}}^{\text{max}}(\mathbf{x}_{\mathbf{n}}) \right)^2 \leq \beta_q(m\kappa)^2 \|U_m^{\text{mod}}\mathbf{X}\|_2^2. \quad (184)$$

Finally, plugging (184) into (180) concludes the proof. \square

A.12 Proof of Proposition 7

Proof. We suppose that Hypothesis 3 is satisfied and we consider $\mathbf{x} \in \mathbb{R}^2$. For a given $n \in \mathbb{N} \setminus \{0\}$, we introduce the random variable

$$\tilde{\mathcal{S}}_{\mathbf{X},n} := \sqrt{\sum_{\|\mathbf{p}\|_\infty \leq n} M_{\mathbf{X}}(\mathbf{x}_{\mathbf{p}})^2}. \quad (185)$$

According to Hypothesis 3, $Z_{\mathbf{X}}(\mathbf{x})$ is jointly independent of $M_{\mathbf{X}}(\mathbf{x}_{\mathbf{p}})$ for $\mathbf{p} \in \{-n, \dots, n\}^2$. Therefore, by composition, $Z_{\mathbf{X}}(\mathbf{x})$ is also independent of $\tilde{\mathcal{S}}_{\mathbf{X},n}$. Moreover, according to (83) and (111), $\tilde{\mathcal{S}}_{\mathbf{X},n}$ converges almost surely towards $\tilde{\mathcal{S}}_{\mathbf{X}}$, which proves independence between $Z_{\mathbf{X}}(\mathbf{x})$ and $\tilde{\mathcal{S}}_{\mathbf{X}}$.

Next, we prove conditional independence between $Z_X(\mathbf{x})$ and $M_X(\mathbf{x})$ given \tilde{S}_X . According to Hypothesis 3,

$$\left(M_X(\mathbf{x}), \tilde{S}_{X,n} \right) \perp\!\!\!\perp Z_X(\mathbf{x}), \quad (186)$$

where $\perp\!\!\!\perp$ stands for independence. This is because $\tilde{S}_{X,n}$ only depends on a finite number of $M_X(\mathbf{x}_p)$. Therefore,

$$Z_X(\mathbf{x}) \perp\!\!\!\perp M_X(\mathbf{x}) \mid \tilde{S}_{X,n}. \quad (187)$$

Finally, since $\tilde{S}_{X,n}$ converges almost surely towards \tilde{S}_X , it comes that $Z_X(\mathbf{x})$ and $M_X(\mathbf{x})$ are conditionally independent given \tilde{S}_X . \square

A.13 Proof of Theorem 2

Proof. We consider $\mathbf{n} \in \mathbb{Z}^2$. By construction, $Q_X(\mathbf{x}_n) := 1 - G_X^{\max}(\mathbf{x}_n)$ only depends on $Z_X(\mathbf{x}_n)$. Therefore, under Hypothesis 3, Proposition 7 implies

$$Q_X(\mathbf{x}_n) \perp\!\!\!\perp M_X(\mathbf{x}_n) \mid \tilde{S}_X^2 \quad \text{and} \quad Q_X(\mathbf{x}_n) \perp\!\!\!\perp \tilde{S}_X^2. \quad (188)$$

Additionally, we introduce

$$\tilde{\Delta}_X := \|\delta_{m,q}X\|_2, \quad (189)$$

where $\delta_{m,q}X$ is defined in (87). Then, using the linearity of \mathbb{E} , we get

$$\begin{aligned} \mathbb{E}[\tilde{\Delta}_X^2 \mid \tilde{S}_X^2 = \sigma] &= \sum_{\mathbf{n} \in \mathbb{Z}^2} \mathbb{E}[\delta_{m,q}[\mathbf{n}]^2 \mid \tilde{S}_X^2 = \sigma] \\ &= \sum_{\mathbf{n} \in \mathbb{Z}^2} \mathbb{E}[U_{m,l}^{\text{mod}}X[\mathbf{n}]^2 (1 - G_X^{\max}(\mathbf{x}_n))^2 \mid \tilde{S}_X^2 = \sigma] \\ &= \sum_{\mathbf{n} \in \mathbb{Z}^2} \mathbb{E}[M_X(\mathbf{x}_n)^2 Q_X(\mathbf{x}_n)^2 \mid \tilde{S}_X^2 = \sigma] \quad (\text{acc. to (83) and (88)}) \\ &= \sum_{\mathbf{n} \in \mathbb{Z}^2} \mathbb{E}[M_X(\mathbf{x}_n)^2 \mid \tilde{S}_X^2 = \sigma] \mathbb{E}[Q_X(\mathbf{x}_n)^2] \quad (\text{acc. to (188)}). \end{aligned}$$

Using the monotone convergence theorem, we get

$$\mathbb{E}[\tilde{\Delta}_X^2 \mid \tilde{S}_X^2 = \sigma] = \mathbb{E}\left[\sum_{\mathbf{n} \in \mathbb{Z}^2} M_X(\mathbf{x}_n)^2 \mid \tilde{S}_X^2 = \sigma \right] \mathbb{E}[Q_X(\mathbf{x}_n)^2]. \quad (190)$$

According to (83) and (111), we have

$$\sum_{\mathbf{n} \in \mathbb{Z}^2} M_X(\mathbf{x}_n)^2 = \|U_m^{\text{mod}}X\|_2^2 = \tilde{S}_X^2. \quad (191)$$

Therefore, we get

$$\begin{aligned} \mathbb{E}[\tilde{\Delta}_X^2 \mid \tilde{S}_X^2 = \sigma] &= \mathbb{E}[\tilde{S}_X^2 \mid \tilde{S}_X^2 = \sigma] \mathbb{E}[Q_X(\mathbf{x}_n)^2] \\ &= \sigma \cdot \mathbb{E}[Q_X(\mathbf{x}_n)^2]. \end{aligned}$$

Under Hypothesis 2, Proposition 6 yields

$$\mathbb{E}[\tilde{\Delta}_X^2 \mid \tilde{S}_X^2 = \sigma] = \sigma \cdot \gamma_q(m\theta)^2. \quad (192)$$

Moreover, we can reformulate \tilde{Q}_X such as defined in (89): $\tilde{Q}_X = \tilde{\Delta}_X/\tilde{S}_X$. Therefore,

$$\mathbb{E}[\tilde{Q}_X^2 \mid \tilde{S}_X^2 = \sigma] = \frac{1}{\sigma} \mathbb{E}[\tilde{\Delta}_X^2 \mid \tilde{S}_X^2 = \sigma] = \gamma_q(m\theta)^2. \quad (193)$$

According to (193), the conditional expected value of \tilde{Q}_X^2 remains the same whatever the outcome of \tilde{S}_X^2 . Thus, the law of total expectation states that

$$\mathbb{E}[\tilde{Q}_X^2] = \mathbb{E}\left[\mathbb{E}[\tilde{Q}_X^2 \mid \tilde{S}_X^2]\right] = \gamma_q(m\boldsymbol{\theta})^2. \quad (194)$$

Since we have assumed Hypothesis 1, we can apply Proposition 5. Using the definition of \tilde{P}_X (86) and \tilde{Q}_X (89), we get

$$\tilde{P}_X \leq \tilde{Q}_X + \beta_q(m\kappa). \quad (195)$$

Then,

$$\mathbb{E}[\tilde{P}_X^2] \leq \mathbb{E}[\tilde{Q}_X^2] + 2\beta_q(m\kappa)\mathbb{E}[\tilde{Q}_X] + \beta_q(m\kappa)^2. \quad (196)$$

According to Jensen's inequality,

$$\mathbb{E}[\tilde{Q}_X] \leq \sqrt{\mathbb{E}[\tilde{Q}_X^2]} = \gamma_q(m\boldsymbol{\theta}). \quad (197)$$

Thus,

$$\mathbb{E}[\tilde{P}_X^2] \leq \gamma_q(m\boldsymbol{\theta})^2 + 2\beta_q(m\kappa)\gamma_q(m\boldsymbol{\theta}) + \beta_q(m\kappa)^2, \quad (198)$$

which yields (112). \square

A.14 Proof of Lemma 6

Proof. First, we show that, for any $\mathbf{x} \in \mathbb{R}^2$,

$$\mathbf{M}_{\mathcal{T}_u X}(\mathbf{x}) = \mathcal{T}_{su} \mathbf{M}_X(\mathbf{x}); \quad (199)$$

$$\mathbf{Z}_{\mathcal{T}_u X}(\mathbf{x}) = \mathcal{T}_{su} \mathbf{Z}_X(\mathbf{x}). \quad (200)$$

According to Lemma 3, and since the convolution product commutes with translations, we have

$$(F_{\mathcal{T}_u X} * \bar{\Psi}_W)(\mathbf{x}) = \mathcal{T}_{su}(F_X * \bar{\Psi}_W)(\mathbf{x}). \quad (201)$$

Then, using (81), the above expression becomes

$$\mathbf{M}_{\mathcal{T}_u X}(\mathbf{x}) \times \mathbf{Z}_{\mathcal{T}_u X}(\mathbf{x}) = (\mathcal{T}_{su} \mathbf{M}_X)(\mathbf{x}) \times (\mathcal{T}_{su} \mathbf{Z}_X)(\mathbf{x}), \quad (202)$$

which yields (199) and (200), by uniqueness of the magnitude-phase decomposition. Finally, we remind that

$$\mathcal{T}_{su} \mathbf{M}_X(\mathbf{x}) = \mathbf{M}_X(\mathbf{x} - su) \quad \text{and} \quad \mathcal{T}_{su} \mathbf{Z}_X(\mathbf{x}) = \mathbf{Z}_X(\mathbf{x} - su). \quad (203)$$

Then, considering hypotheses Hypotheses 2 and 3 with $\mathbf{x} \leftarrow \mathbf{x} - su$ yields the result. \square

A.15 Proof of Theorem 3

Proof. Using the triangle inequality, we compute

$$\begin{aligned} & \|U_{m,q}^{\max}(\mathcal{T}_u X) - U_{m,q}^{\max} X\|_2 \\ & \leq \|U_m^{\text{mod}}(\mathcal{T}_u X)\|_2 \tilde{P}_{\mathcal{T}_u X} + \|U_m^{\text{mod}} X\|_2 \tilde{P}_X + \|U_m^{\text{mod}}(\mathcal{T}_u X) - U_m^{\text{mod}} X\|_2, \end{aligned} \quad (204)$$

where \tilde{P}_X and $\tilde{P}_{\mathcal{T}_u X}$ are defined in (86). According to (38), we can apply Proposition 3 on the first term of (204):

$$\|U_m^{\text{mod}}(\mathcal{T}_u X)\|_2 = \|U_m^{\text{mod}} X\|_2. \quad (205)$$

Moreover, we can apply Theorem 1 to the third term of (204):

$$\|U_m^{\text{mod}}(\mathcal{T}_u X) - U_m^{\text{mod}} X\|_2 \leq \alpha(\kappa u) \|U_m^{\text{mod}} X\|_2. \quad (206)$$

We therefore get

$$\|U_{m,q}^{\max}(\mathcal{T}_u X) - U_{m,q}^{\max} X\|_2 \leq [\tilde{P}_{\mathcal{T}_u X} + \tilde{P}_X + \alpha(\kappa u)] \|U_m^{\text{mod}} X\|_2. \quad (207)$$

Then, by linearity of \mathbb{E} , we get

$$\mathbb{E}[\tilde{\mathbf{R}}_{X, \mathbf{u}}] \leq \mathbb{E}[\tilde{\mathbf{P}}_{\mathcal{T}_{\mathbf{u}}X}] + \mathbb{E}[\tilde{\mathbf{P}}_X] + \alpha(\kappa\mathbf{u}), \quad (208)$$

where $\tilde{\mathbf{R}}_{X, \mathbf{u}}$ has been introduced in (113).

For any stochastic process X' satisfying Hypotheses 2 and 3, Theorem 2 and Jensen's inequality yield:

$$\mathbb{E}[\tilde{\mathbf{P}}_{X'}] \leq \beta_q(m\kappa) + \gamma_q(m\boldsymbol{\theta}). \quad (209)$$

According to Lemma 6, Hypotheses 2 and 3 are also satisfied for $X \leftarrow \mathcal{T}_{\mathbf{u}}X$. Therefore, (209) is valid for both $X' \leftarrow X$ and $X' \leftarrow \mathcal{T}_{\mathbf{u}}X$, and plugging it into (208) concludes the proof. \square

B Appendix – Theoretical Foundations for our Hypotheses

In this section, we provide theoretical arguments for justifying Hypotheses 2 and 3. As will be discussed, these hypotheses rely on some degree of shift invariance for input images, which implies the notions of stationarity and phase-shift-equivariance for stochastic processes.

B.1 Stationary and Local Stationarity

Given $n \in \mathbb{N} \setminus \{0\}$, we define *n-th order stationarity* of a given stochastic process \mathbf{F} as stated by Park & Park (2018, p. 152): For any $n' \leq n$, $(\mathbf{x}_0, \dots, \mathbf{x}_{n'-1}) \in (\mathbb{R}^2)^{n'}$ and $\mathbf{h} \in \mathbb{R}^2$, the joint distribution of $(\mathbf{F}(\mathbf{x}_0), \dots, \mathbf{F}(\mathbf{x}_{n'-1}))$ is identical to the one of $(\mathcal{T}_{\mathbf{h}}\mathbf{F}(\mathbf{x}_0), \dots, \mathcal{T}_{\mathbf{h}}\mathbf{F}(\mathbf{x}_{n'-1}))$. Additionally, *strict-sense stationarity* is defined as *n-th order stationarity* for any $n \in \mathbb{N} \setminus \{0\}$.

Strict-sense stationarity suggests that any translated version of a given image is equally likely. However, this property is seldom fully achieved for images (Tygert et al., 2016). First, \mathbf{X} has fixed boundaries. Consequently, any realization of $\mathbf{F}_X(\mathbf{x})$ quickly vanishes as $\|\mathbf{x}\|$ tends to ∞ . Furthermore, depending on which category the image belongs to, the pixel distribution is likely to vary across various regions. For instance, we can expect the main subject to be located at the center of the image. We refer readers to Torralba & Oliva (2003) for more details on statistical properties of images from natural versus man-made objects. For this reason, we introduce a notion of local stationarity as follows.

Definition 1. Given $\tau \geq 0$, a stochastic process \mathbf{F} is *τ -locally stationary* to the *n-th order* if, for any $0 < n'' \leq n' \leq n$, any $(\mathbf{x}_0, \dots, \mathbf{x}_{n'-1}) \in (\mathbb{R}^2)^{n'}$, any displacement vector $\mathbf{h} \in \mathbb{R}^2$, and any pair of measurable sets $\mathfrak{X}, \mathfrak{Y} \subset \mathbb{R}^{n''} \times \mathbb{R}^{n'-n''}$,

$$\left| \mathbb{P}\left\{ \left(\mathcal{T}_{\mathbf{h}}\mathbf{F}(\mathbf{x}_0), \dots, \mathcal{T}_{\mathbf{h}}\mathbf{F}(\mathbf{x}_{n''-1}) \right)^\top \in \mathfrak{X} \mid \left(\mathcal{T}_{\mathbf{h}}\mathbf{F}(\mathbf{x}_{n''}), \dots, \mathcal{T}_{\mathbf{h}}\mathbf{F}(\mathbf{x}_{n'-1}) \right)^\top \in \mathfrak{Y} \right\} - \mathbb{P}\left\{ \left(\mathbf{F}(\mathbf{x}_0), \dots, \mathbf{F}(\mathbf{x}_{n''-1}) \right)^\top \in \mathfrak{X} \mid \left(\mathbf{F}(\mathbf{x}_{n''}), \dots, \mathbf{F}(\mathbf{x}_{n'-1}) \right)^\top \in \mathfrak{Y} \right\} \right| \leq \tau \|\mathbf{h}\|_2. \quad (210)$$

Moreover, \mathbf{F} is *strict-sense τ -locally stationary* if this property is satisfied for any $n \in \mathbb{N}$.

Remark 10. The special case where $\tau = 0$ corresponds to the standard definition of *n-th order stationarity*. Additionally, the use of conditional probabilities is essential for defining local stationarity, as it helps prevent instabilities that may arise when conditioning on low-probability events.

In what follows, we consider the following narrowband stochastic process:

$$\mathbf{F}_{1, X} : \mathbf{x} \mapsto (\mathbf{F}_X * \bar{\Psi}_W)(\mathbf{x}). \quad (211)$$

We assume that \mathbf{F}_X , and therefore $\mathbf{F}_{1, X}$, is nearly shift-invariant for displacement vectors that are much smaller than the image “characteristic” size in the continuous domain, which is equal to sN , where, as a reminder, $N \in \mathbb{N}$ denotes the support size of input images and $s > 0$ denotes the sampling interval. More formally, $\mathbf{F}_{1, X}$ is assumed to be strict-sense τ -locally stationary with

$$\tau := \frac{1}{sN}. \quad (212)$$

B.2 Translations and Phase Shifts

We consider the following stochastic processes:

$$F_{0,X} : \mathbf{x} \mapsto (F_X * \bar{\Psi}_W)(\mathbf{x}) e^{i\langle \boldsymbol{\nu}, \mathbf{x} \rangle} \quad \text{and} \quad F_{1,X} : \mathbf{x} \mapsto (F_X * \bar{\Psi}_W)(\mathbf{x}). \quad (213)$$

To justify our hypotheses, we also need to characterize $\mathcal{T}_h F_{1,X}$ as a function of $F_{1,X}$. Having:

$$\mathcal{T}_h F_{1,X}(\mathbf{x}) = \mathcal{T}_h F_{0,X}(\mathbf{x}) e^{-i\langle \boldsymbol{\nu}, \mathbf{x} \rangle} e^{i\langle \boldsymbol{\nu}, \mathbf{h} \rangle}, \quad (214)$$

we get

$$|\mathcal{T}_h F_{1,X}(\mathbf{x}) - F_{1,X}(\mathbf{x}) e^{i\langle \boldsymbol{\nu}, \mathbf{h} \rangle}| = |\mathcal{T}_h F_{0,X}(\mathbf{x}) - F_{0,X}(\mathbf{x})|. \quad (215)$$

According to Lemma 1, the support of $\widehat{F}_{0,X}$ is contained within the ball $B_\infty(\frac{\kappa}{2s})$. Intuitively, we can define a ‘‘minimal wavelength’’ $\lambda := 2\pi s/\kappa$ such that, if $\|\mathbf{h}\|_2 \ll \lambda$, we can approximate $\mathcal{T}_h F_{0,X}(\mathbf{x}) \approx F_{0,X}(\mathbf{x})$, and therefore

$$\mathcal{T}_h F_{1,X}(\mathbf{x}) \approx F_{1,X}(\mathbf{x}) e^{i\langle \boldsymbol{\nu}, \mathbf{h} \rangle}. \quad (216)$$

If the two terms were strictly identical for any $\mathbf{x} \in \mathbb{R}^2$, we would get, for any $n \in \mathbb{N}$, $(\mathbf{x}_0, \dots, \mathbf{x}_n) \in (\mathbb{R}^2)^n$, and any measurable set $\mathfrak{X} \subset \mathbb{R}^n$,

$$(F_{1,X}(\mathbf{x}_0), \dots, F_{1,X}(\mathbf{x}_{n-1}))^\top \in \mathfrak{X} \iff (\mathcal{T}_h F_{1,X}(\mathbf{x}_0), \dots, \mathcal{T}_h F_{1,X}(\mathbf{x}_{n-1}))^\top \in e^{i\langle \boldsymbol{\nu}, \mathbf{h} \rangle} \mathfrak{X}, \quad (217)$$

and therefore,

$$\mathbb{P} \left\{ (F_{1,X}(\mathbf{x}_0), \dots, F_{1,X}(\mathbf{x}_{n-1}))^\top \in \mathfrak{X} \right\} = \mathbb{P} \left\{ (\mathcal{T}_h F_{1,X}(\mathbf{x}_0), \dots, \mathcal{T}_h F_{1,X}(\mathbf{x}_{n-1}))^\top \in e^{i\langle \boldsymbol{\nu}, \mathbf{h} \rangle} \mathfrak{X} \right\}. \quad (218)$$

Instead, we relax the above equality (218) and introduce the following definition.

Definition 2. Given $\tau \geq 0$, a stochastic process F is τ -locally phase-shift-equivariant to the n -th order with respect to the frequency vector $\boldsymbol{\nu} \in \mathbb{R}^2$ if, for any $0 < n'' \leq n' \leq n$, any $(\mathbf{x}_0, \dots, \mathbf{x}_{n'-1}) \in (\mathbb{R}^2)^{n'}$, any displacement vector $\mathbf{h} \in \mathbb{R}^2$, and any pair of measurable sets $\mathfrak{X}, \mathfrak{Y} \subset \mathbb{R}^{n''} \times \mathbb{R}^{n'-n''}$,

$$\left| \mathbb{P} \left\{ (\mathcal{T}_h F(\mathbf{x}_0), \dots, \mathcal{T}_h F(\mathbf{x}_{n''-1}))^\top \in e^{i\langle \boldsymbol{\nu}, \mathbf{h} \rangle} \mathfrak{X} \mid (\mathcal{T}_h F(\mathbf{x}_{n''}), \dots, \mathcal{T}_h F(\mathbf{x}_{n'-1}))^\top \in e^{i\langle \boldsymbol{\nu}, \mathbf{h} \rangle} \mathfrak{Y} \right\} - \mathbb{P} \left\{ (F(\mathbf{x}_0), \dots, F(\mathbf{x}_{n''-1}))^\top \in \mathfrak{X} \mid (F(\mathbf{x}_{n''}), \dots, F(\mathbf{x}_{n'-1}))^\top \in \mathfrak{Y} \right\} \right| \leq \tau \|\mathbf{h}\|_2. \quad (219)$$

Moreover, F is *strict-sense* τ -locally phase-shift-equivariant if this property is satisfied for any $n \in \mathbb{N}$.

In what follows, we will assume that the stochastic process $F_{1,X}$ is nearly phase-shift-equivariant for displacement vectors that are much smaller than the minimal wavelength λ . More formally, $F_{1,X}$ is assumed to be strict-sense τ' -locally phase-shift-equivariant with

$$\tau' := 1/\lambda = \frac{\kappa}{2\pi s}. \quad (220)$$

B.3 Justification for Hypothesis 2

We then consider the following proposition, which states that the probability measure of $Z_X(\mathbf{x})$ is nearly invariant with respect to phase shifts.

Proposition 8. *We assume that $F_{1,X}$ is τ -locally stationary (Definition 1) and τ' -locally phase-shift-equivariant with respect to $\boldsymbol{\nu}$ (Definition 2), both to the first order. Then, for any measurable set $\mathfrak{A} \subset \mathbb{S}^1$,*

$$\forall \omega \in [0, 2\pi], \left| \mu(\mathfrak{A}) - \mu(e^{i\omega} \mathfrak{A}) \right| \leq 2\pi(\tau + \tau') / \|\boldsymbol{\nu}\|_2. \quad (221)$$

where $\mu : \mathfrak{A} \mapsto \mathbb{P} \{ Z_X(\mathbf{x}) \in \mathfrak{A} \}$ denotes the probability measure of $Z_X(\mathbf{x})$.

Proof. Let $\omega \in [0, 2\pi]$. We compute:

$$\begin{aligned} |\mu(\mathfrak{A}) - \mu(e^{i\omega}\mathfrak{A})| &= |\mathbb{P}\{Z_X(\mathbf{x}) \in \mathfrak{A}\} - \mathbb{P}\{Z_X(\mathbf{x}) \in e^{i\omega}\mathfrak{A}\}| \\ &= |\mathbb{P}\{Z_X(\mathbf{x}) \in \mathfrak{A}\} - \mathbb{P}\{Z_X(\mathbf{x}) \in e^{i\langle \boldsymbol{\nu}, \mathbf{h} \rangle}\mathfrak{A}\}|, \end{aligned}$$

where we have denoted $\mathbf{h} := \omega \boldsymbol{\nu} / \|\boldsymbol{\nu}\|_2^2$. Then, using the triangle inequality,

$$\begin{aligned} |\mu(\mathfrak{A}) - \mu(e^{i\omega}\mathfrak{A})| &\leq |\mathbb{P}\{Z_X(\mathbf{x}) \in \mathfrak{A}\} - \mathbb{P}\{\mathcal{T}_{\mathbf{h}}Z_X(\mathbf{x}) \in e^{i\langle \boldsymbol{\nu}, \mathbf{h} \rangle}\mathfrak{A}\}| \\ &\quad + |\mathbb{P}\{Z_X(\mathbf{x}) \in e^{i\langle \boldsymbol{\nu}, \mathbf{h} \rangle}\mathfrak{A}\} - \mathbb{P}\{\mathcal{T}_{\mathbf{h}}Z_X(\mathbf{x}) \in e^{i\langle \boldsymbol{\nu}, \mathbf{h} \rangle}\mathfrak{A}\}| \\ &= |\mathbb{P}\{F_{1,X}(\mathbf{x}) \in u^{-1}(\mathfrak{A})\} - \mathbb{P}\{\mathcal{T}_{\mathbf{h}}F_{1,X}(\mathbf{x}) \in u^{-1}(e^{i\langle \boldsymbol{\nu}, \mathbf{h} \rangle}\mathfrak{A})\}| \\ &\quad + |\mathbb{P}\{F_{1,X}(\mathbf{x}) \in u^{-1}(e^{i\langle \boldsymbol{\nu}, \mathbf{h} \rangle}\mathfrak{A})\} - \mathbb{P}\{\mathcal{T}_{\mathbf{h}}F_{1,X}(\mathbf{x}) \in u^{-1}(e^{i\langle \boldsymbol{\nu}, \mathbf{h} \rangle}\mathfrak{A})\}|, \end{aligned}$$

where we have denoted $u : z \mapsto e^{i\angle(z)}$. We can easily show that $u^{-1}(e^{i\langle \boldsymbol{\nu}, \mathbf{h} \rangle}\mathfrak{A}) = e^{i\langle \boldsymbol{\nu}, \mathbf{h} \rangle}u^{-1}(\mathfrak{A})$. Therefore, the above expression can be re-written:

$$\begin{aligned} |\mu(\mathfrak{A}) - \mu(e^{i\omega}\mathfrak{A})| &\leq |\mathbb{P}\{F_{1,X}(\mathbf{x}) \in \mathfrak{X}\} - \mathbb{P}\{\mathcal{T}_{\mathbf{h}}F_{1,X}(\mathbf{x}) \in e^{i\langle \boldsymbol{\nu}, \mathbf{h} \rangle}\mathfrak{X}\}| \\ &\quad + |\mathbb{P}\{F_{1,X}(\mathbf{x}) \in \mathfrak{X}'\} - \mathbb{P}\{\mathcal{T}_{\mathbf{h}}F_{1,X}(\mathbf{x}) \in \mathfrak{X}'\}|, \end{aligned} \quad (222)$$

where we have denoted $\mathfrak{X} := u^{-1}(\mathfrak{A})$ and $\mathfrak{X}' := u^{-1}(e^{i\langle \boldsymbol{\nu}, \mathbf{h} \rangle}\mathfrak{A})$. Then, by applying (210) and (219) with $n' = 1$ to (222), we get:

$$|\mu(\mathfrak{A}) - \mu(e^{i\omega}\mathfrak{A})| \leq (\tau + \tau') \|\mathbf{h}\|_2. \quad (223)$$

Finally, using the definition of \mathbf{h} and bounding ω by 2π yields the result. \square

If we replace τ and τ' in (221) by their respective values in (212) and (220), we get,

$$\forall \omega \in [0, 2\pi], |\mu(\mathfrak{A}) - \mu(e^{i\omega}\mathfrak{A})| \leq \frac{2\pi}{sN} \times \|\boldsymbol{\nu}\|_2^{-1} + \frac{\kappa}{s} \times \|\boldsymbol{\nu}\|_2^{-1}. \quad (224)$$

We recall that $\boldsymbol{\nu} := \boldsymbol{\theta}/s$. Then, by applying the constraints stated in (84) and (85), we get

$$\forall \omega \in [0, 2\pi], |\mu(\mathfrak{A}) - \mu(e^{i\omega}\mathfrak{A})| \ll 1. \quad (225)$$

Therefore, μ is almost invariant to phase shifts. Since the only probability measure satisfying the phase-shift invariance property is the uniform probability measure,³ we deduce that $Z_X(\mathbf{x})$ follows a near-uniform distribution on \mathbb{S}^1 . For the sake of simplicity, in Hypothesis 2 we have assumed a strictly-uniform distribution.

B.4 Justification for Hypothesis 3

Let $n \in \mathbb{N} \setminus \{0\}$ and $\mathbf{x}, \mathbf{y}_0, \dots, \mathbf{y}_{n-1} \in \mathbb{R}^2$. To simplify notations, we consider the random vector

$$\mathbf{M} := (\mathbf{M}_X(\mathbf{y}_0), \dots, \mathbf{M}_X(\mathbf{y}_{n-1}))^\top \quad (226)$$

with outcomes in \mathbb{R}_+^n . This section is organized as follows. Using reasoning similar to that in Proposition 8, we show that, for any measurable subset $\mathfrak{S} \subset \mathbb{R}_+^n$, Z_X follows a near-uniform probability distribution conditionally to $\mathbf{M} \in \mathfrak{S}$. Since we already assumed that Z_X follows an unconditional uniform distribution, we deduce that Z_X and \mathbf{M} are nearly independent.

³Any probability measure defined on \mathbb{S}^1 is a Radon measure. Therefore, according to Haar's theorem (Halmos, 2013), there exists a unique probability measure on \mathbb{S}^1 satisfying the phase-shift invariance property, and it turns out that the uniform probability measure is one such candidate.

Proposition 9. *We assume that $F_{1,X}$ is τ -locally stationary (Definition 1) and τ' -locally phase-shift-equivariant with respect to ν (Definition 2), both in the strict sense. Then, for any measurable sets $\mathfrak{A} \subset \mathbb{S}^1$ and $\mathfrak{S} \subset \mathbb{R}_+^n$,*

$$\forall \omega \in [0, 2\pi], \quad |\mu_{\mathfrak{S}}(\mathfrak{A}) - \mu_{\mathfrak{S}}(e^{i\omega}\mathfrak{A})| \leq 2\pi(\tau + \tau') / \|\nu\|_2. \quad (227)$$

where $\mu_{\mathfrak{S}} : \mathfrak{A} \mapsto \mathbb{P}\{Z_X(\mathbf{x}) \in \mathfrak{A} \mid \mathbf{M} \in \mathfrak{S}\}$ denotes the conditional probability measure of $Z_X(\mathbf{x})$.

Proof. Let $\omega \in [0, 2\pi]$. We compute:

$$\begin{aligned} |\mu_{\mathfrak{S}}(e^{i\omega}\mathfrak{A}) - \mu_{\mathfrak{S}}(\mathfrak{A})| &= \left| \mathbb{P}\{Z_X(\mathbf{x}) \in e^{i\omega}\mathfrak{A} \mid \mathbf{M} \in \mathfrak{S}\} - \mathbb{P}\{Z_X(\mathbf{x}) \in \mathfrak{A} \mid \mathbf{M} \in \mathfrak{S}\} \right| \\ &= \left| \mathbb{P}\{Z_X(\mathbf{x}) \in e^{i\langle \nu, \mathbf{h} \rangle} \mathfrak{A} \mid \mathbf{M} \in \mathfrak{S}\} - \mathbb{P}\{Z_X(\mathbf{x}) \in \mathfrak{A} \mid \mathbf{M} \in \mathfrak{S}\} \right|, \end{aligned}$$

where we have denoted

$$\mathbf{h} := \omega \nu / \|\nu\|_2^2 \quad (228)$$

Then, using the triangle inequality,

$$\begin{aligned} |\mu_{\mathfrak{S}}(e^{i\omega}\mathfrak{A}) - \mu_{\mathfrak{S}}(\mathfrak{A})| &\leq \left| \mathbb{P}\{\mathcal{T}_{\mathbf{h}}Z_X(\mathbf{x}) \in e^{i\langle \nu, \mathbf{h} \rangle} \mathfrak{A} \mid \mathcal{T}_{\mathbf{h}}\mathbf{M} \in \mathfrak{S}\} - \mathbb{P}\{Z_X(\mathbf{x}) \in e^{i\langle \nu, \mathbf{h} \rangle} \mathfrak{A} \mid \mathbf{M} \in \mathfrak{S}\} \right| \\ &\quad + \left| \mathbb{P}\{\mathcal{T}_{\mathbf{h}}Z_X(\mathbf{x}) \in e^{i\langle \nu, \mathbf{h} \rangle} \mathfrak{A} \mid \mathcal{T}_{\mathbf{h}}\mathbf{M} \in \mathfrak{S}\} - \mathbb{P}\{Z_X(\mathbf{x}) \in \mathfrak{A} \mid \mathbf{M} \in \mathfrak{S}\} \right|, \end{aligned} \quad (229)$$

where we have denoted

$$\mathcal{T}_{\mathbf{h}}\mathbf{M} := (\mathcal{T}_{\mathbf{h}}M_X(\mathbf{y}_0), \dots, \mathcal{T}_{\mathbf{h}}M_X(\mathbf{y}_{n-1}))^\top. \quad (230)$$

Let us split this expression for the sake of readability. The first term after the \leq sign can be equivalently written as follows:

$$\begin{aligned} &\left| \mathbb{P}\{\mathcal{T}_{\mathbf{h}}Z_X(\mathbf{x}) \in e^{i\langle \nu, \mathbf{h} \rangle} \mathfrak{A} \mid \mathcal{T}_{\mathbf{h}}\mathbf{M} \in \mathfrak{S}\} - \mathbb{P}\{Z_X(\mathbf{x}) \in e^{i\langle \nu, \mathbf{h} \rangle} \mathfrak{A} \mid \mathbf{M} \in \mathfrak{S}\} \right| \\ &= \left| \mathbb{P}\{\mathcal{T}_{\mathbf{h}}F_{1,X}(\mathbf{x}) \in \mathfrak{X}' \mid (\mathcal{T}_{\mathbf{h}}F_{1,X}(\mathbf{y}_0), \dots, \mathcal{T}_{\mathbf{h}}F_{1,X}(\mathbf{y}_{n-1}))^\top \in \mathfrak{Y}\} \right. \\ &\quad \left. - \mathbb{P}\{F_{1,X}(\mathbf{x}) \in \mathfrak{X}' \mid (F_{1,X}(\mathbf{y}_0), \dots, F_{1,X}(\mathbf{y}_{n-1}))^\top \in \mathfrak{Y}\} \right|, \end{aligned} \quad (231)$$

where we have denoted

$$\mathfrak{X}' := u^{-1}(e^{i\langle \nu, \mathbf{h} \rangle} \mathfrak{A}) \quad \text{and} \quad \mathfrak{Y} := \mathbf{v}^{-1}(\mathfrak{S}), \quad (232)$$

with

$$u : z \mapsto e^{i\angle(z)} \quad \text{and} \quad \mathbf{v} : (z_0, \dots, z_{n-1})^\top \mapsto (|z_0|, \dots, |z_{n-1}|)^\top. \quad (233)$$

Therefore, according the local stationarity hypothesis, we apply (210) to (231), which yields

$$\left| \mathbb{P}\{\mathcal{T}_{\mathbf{h}}Z_X(\mathbf{x}) \in e^{i\langle \nu, \mathbf{h} \rangle} \mathfrak{A} \mid \mathcal{T}_{\mathbf{h}}\mathbf{M} \in \mathfrak{S}\} - \mathbb{P}\{Z_X(\mathbf{x}) \in e^{i\langle \nu, \mathbf{h} \rangle} \mathfrak{A} \mid \mathbf{M} \in \mathfrak{S}\} \right| \leq \tau \|\mathbf{h}\|_2. \quad (234)$$

Next, the second term after the \leq sign in (229) can be equivalently written as follows:

$$\begin{aligned} &\left| \mathbb{P}\{\mathcal{T}_{\mathbf{h}}Z_X(\mathbf{x}) \in e^{i\langle \nu, \mathbf{h} \rangle} \mathfrak{A} \mid \mathcal{T}_{\mathbf{h}}\mathbf{M} \in \mathfrak{S}\} - \mathbb{P}\{Z_X(\mathbf{x}) \in \mathfrak{A} \mid \mathbf{M} \in \mathfrak{S}\} \right| \\ &= \left| \mathbb{P}\{\mathcal{T}_{\mathbf{h}}F_{1,X}(\mathbf{x}) \in u^{-1}(e^{i\langle \nu, \mathbf{h} \rangle} \mathfrak{A}) \mid (\mathcal{T}_{\mathbf{h}}F_{1,X}(\mathbf{y}_0), \dots, \mathcal{T}_{\mathbf{h}}F_{1,X}(\mathbf{y}_{n-1}))^\top \in \mathbf{v}^{-1}(\mathfrak{S})\} \right. \\ &\quad \left. - \mathbb{P}\{F_{1,X}(\mathbf{x}) \in u^{-1}(\mathfrak{A}) \mid (F_{1,X}(\mathbf{y}_0), \dots, F_{1,X}(\mathbf{y}_{n-1}))^\top \in \mathbf{v}^{-1}(\mathfrak{S})\} \right|. \end{aligned} \quad (235)$$

As explained in the proof of Proposition 8, we have

$$u^{-1}(e^{i\langle \nu, \mathbf{h} \rangle} \mathfrak{A}) = e^{i\langle \nu, \mathbf{h} \rangle} u^{-1}(\mathfrak{A}). \quad (236)$$

Moreover, we can easily show that

$$\mathbf{v}^{-1}(\mathfrak{S}) = e^{i\langle \nu, \mathbf{h} \rangle} \mathbf{v}^{-1}(\mathfrak{S}). \quad (237)$$

Therefore, (235) can be re-written:

$$\begin{aligned} & \left| \mathbb{P}\{\mathcal{T}_{\mathbf{h}}Z_X(\mathbf{x}) \in e^{i\langle \nu, \mathbf{h} \rangle} \mathfrak{A} \mid \mathcal{T}_{\mathbf{h}}\mathbf{M} \in \mathfrak{S}\} - \mathbb{P}\{Z_X(\mathbf{x}) \in \mathfrak{A} \mid \mathbf{M} \in \mathfrak{S}\} \right| \\ &= \left| \mathbb{P}\{\mathcal{T}_{\mathbf{h}}F_{1,X}(\mathbf{x}) \in e^{i\langle \nu, \mathbf{h} \rangle} \mathfrak{X} \mid (\mathcal{T}_{\mathbf{h}}F_{1,X}(\mathbf{y}_0), \dots, \mathcal{T}_{\mathbf{h}}F_{1,X}(\mathbf{y}_{n-1}))^\top \in e^{i\langle \nu, \mathbf{h} \rangle} \mathfrak{Y}\} \right. \\ & \quad \left. - \mathbb{P}\{F_{1,X}(\mathbf{x}) \in \mathfrak{X} \mid (F_{1,X}(\mathbf{y}_0), \dots, F_{1,X}(\mathbf{y}_{n-1}))^\top \in \mathfrak{Y}\} \right|, \end{aligned} \quad (238)$$

where we have denoted

$$\mathfrak{X} := u^{-1}(\mathfrak{A}) \quad \text{and} \quad \mathfrak{Y} := \mathbf{v}^{-1}(\mathfrak{S}). \quad (239)$$

Therefore, according the local phase-shift equivariance hypothesis, we apply (219) to (238), which yields

$$\left| \mathbb{P}\{\mathcal{T}_{\mathbf{h}}Z_X(\mathbf{x}) \in e^{i\langle \nu, \mathbf{h} \rangle} \mathfrak{A} \mid \mathcal{T}_{\mathbf{h}}\mathbf{M} \in \mathfrak{S}\} - \mathbb{P}\{Z_X(\mathbf{x}) \in \mathfrak{A} \mid \mathbf{M} \in \mathfrak{S}\} \right| \leq \tau' \|\mathbf{h}\|_2. \quad (240)$$

Finally, plugging (234) and (240) into (229), using the definition of \mathbf{h} in (228), and bounding ω by 2π , yields the result. \square

By applying the constraints sated in (84) and (85), we get (see Appendix B.3),

$$\forall \omega \in [0, 2\pi], \quad |\mu_{\mathfrak{S}}(\mathfrak{A}) - \mu_{\mathfrak{S}}(e^{i\omega} \mathfrak{A})| \ll 1. \quad (241)$$

Therefore, $Z_X(\mathbf{x})$ follows a near-uniform conditional distribution on \mathbb{S}^1 given $\mathbf{M} \in \mathfrak{S}$.

Furthermore, according to Appendix B.3, $Z_X(\mathbf{x})$ also follows a near-uniform unconditional distribution. Therefore, we get, for any measurable sets $\mathfrak{A} \subset \mathbb{S}^1$ and $\mathfrak{S} \subset \mathbb{R}_+^n$,

$$\mathbb{P}\{Z_X(\mathbf{x}) \in \mathfrak{A} \mid (\mathbf{M} \in \mathfrak{S})\} \approx \mathbb{P}\{Z_X(\mathbf{x}) \in \mathfrak{A}\}. \quad (242)$$

We deduce that $Z_X(\mathbf{x})$ and \mathbf{M} are nearly independent. For the sake of simplicity, in Hypothesis 3 we have assumed strict independence.

C Appendix – Details on DT-CWPT

A description of the transform itself is provided in Appendix C.1. Then, Appendix C.2 shows that DT-CWPT performs convolutions with a subsampling factor m_J which depends on the decomposition depth J . Finally, the Gabor-like nature of the convolution kernels is established in Appendix C.3.

C.1 Background

We provide a brief overview of the classical, real-valued 2D wavelet packet transform (WPT) algorithm (Mallat, 2009, p. 377), before introducing the redundant, complex-valued and oriented DT-CWPT (Bayram & Selesnick, 2008).

C.1.1 Discrete Wavelet Packet Transform

Given a pair of low- and high-pass 1D orthogonal filters $h, g \in l_{\mathbb{R}}^2(\mathbb{Z})$ satisfying a *quadrature mirror filter* (QMF) relationship, we consider a separable 2D filter bank (FB), denoted by $\mathbf{G} := (G_l)_{l \in \{0, \dots, 3\}}$, defined by

$$G_0 = h \otimes h; \quad G_1 = h \otimes g; \quad G_2 = g \otimes h; \quad G_3 = g \otimes g. \quad (243)$$

Let $X \in l_{\mathbb{R}}^2(\mathbb{Z})$. The decomposition starts with $D_0^{(0)} = X$. Given $j \in \mathbb{N}$, suppose that we have computed 4^j sequences of wavelet packet coefficients at stage j , denoted by $D_l^{(j)} \in l_{\mathbb{R}}^2(\mathbb{Z})$ for each $l \in \{0, \dots, 4^j - 1\}$. They are referred to as *feature maps*.

At stage $j + 1$, we compute a new representation of X with increased frequency resolution—and decreased spatial resolution. It is obtained by further decomposing each feature map $D_l^{(j)}$ into four sub-sequences, using subsampled (or strided) convolutions with kernels G_k , for each $k \in \{0, \dots, 3\}$:

$$\forall k \in \{0, \dots, 3\}, D_{4l+k}^{(j+1)} = (D_l^{(j)} * \overline{G_k}) \downarrow 2. \quad (244)$$

The algorithm stops after reaching the desired number of stages $J > 0$ —referred to as *decomposition depth*. Then,

$$\mathbf{D}^{(J)} := (D_l^{(J)})_{l \in \{0, \dots, 4^J - 1\}} \quad (245)$$

constitutes a multichannel representation of X in an orthonormal basis, from which the original image can be retrieved.

C.1.2 Dual-Tree Complex Wavelet Packet Transform

Despite having interesting properties such as sparse signal representation, WPT is unstable with respect to small shifts and suffers from a poor directional selectivity. To overcome this, Kingsbury (2001) designed a new type of discrete wavelet transform, where images are decomposed in a redundant frame of nearly-analytic, complex-valued waveforms. It was later extended to the wavelet packet framework by Bayram & Selesnick (2008). The latter operation, referred to as *dual-tree complex wavelet packet transform* (DT-CWPT), is performed as follows.

Let $(h^{[0]}, g^{[0]})$ and $(h^{[1]}, g^{[1]})$ denote two pairs of QMFs as defined in Appendix C.1.1, satisfying the *half-sample delay* condition:

$$\forall \omega \in [-\pi, \pi], \widehat{h}^{[1]}(\omega) = e^{-i\omega/2} \widehat{h}^{[0]}(\omega). \quad (246)$$

Then, for any $k \in \{0, \dots, 3\}$, we build a 2D FB $\mathbf{G}_k := (G_{k,l})_{l \in \{0, \dots, 3\}}$ similarly to (243):

$$\mathbf{G}_{k,0} = h_i \otimes h_j; \quad \mathbf{G}_{k,1} = h_i \otimes g_j; \quad \mathbf{G}_{k,2} = g_i \otimes h_j; \quad \mathbf{G}_{k,3} = g_i \otimes g_j, \quad (247)$$

where $i, j \in \{0, 1\}$ are defined such that $k = 2 \times i + j$.⁴

Let $J > 0$ denote a decomposition depth. Using each of the four FBs \mathbf{G}_{0-3} as defined above, we assume that we have decomposed an input image X into four multichannel WPT representations $\mathbf{D}_{0-3}^{(J)}$, each of which satisfies (244) and (245). Then, for any $l \in \{0, \dots, 4^J - 1\}$, the following complex feature maps are computed:

$$\begin{pmatrix} D_l^{\nearrow(J)} \\ D_l^{\searrow(J)} \end{pmatrix} = \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} D_l^{[0](J)} \\ D_l^{[3](J)} \end{pmatrix} - i \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} D_l^{[2](J)} \\ D_l^{[1](J)} \end{pmatrix}. \quad (248)$$

As explained in Appendix C.3, the feature maps of dual-tree coefficients have their Fourier transform restricted to a compact region of the frequency plane, and as such can be considered as Gabor-like coefficients. In the above expression, the arrow points to the Fourier quadrant where energy is concentrated. Furthermore, in the specific case where input images are real-valued, $D_l^{\swarrow(J)}$ and $D_l^{\nwarrow(J)}$ are defined as the complex conjugates of the above feature maps, and therefore do not need to be explicitly computed. Then,

$$\mathbf{D}^{(J)} := (D_l^{\nearrow(J)}, D_l^{\searrow(J)}, D_l^{\swarrow(J)}, D_l^{\nwarrow(J)})_{l \in \{0, \dots, 4^J - 1\}} \quad (249)$$

constitutes a complex-valued, four-time redundant multichannel representation of X from which the original image can be reconstructed.

⁴Actually, the FB design requires some technicalities which are not described here.

C.2 Convolution Operators

We now show that DT-CWPT performs subsampled convolutions with Gabor-like filters, whose characteristics will be specified. First, we state the following lemma concerning the real-valued WPT algorithm, such as introduced in Appendix C.1.1. It is a simple reformulation of the well-known result that two successive convolutions can be written as another convolution with a wider kernel.

Lemma 7. *For any $l \in \{0, \dots, 4^J - 1\}$, there exists $V_l^{(J)} \in l_{\mathbb{R}}^2(\mathbb{Z}^2)$ such that*

$$D_l^{(J)} = (X * \bar{V}_l^{(J)}) \downarrow 2^J. \quad (250)$$

Proof. We introduce the upsampling operator: $(X \uparrow m)[\mathbf{n}] := X[\mathbf{n}/m]$ if $\mathbf{n}/m \in \mathbb{Z}^2$, and 0 otherwise. We also consider the “identity” filter $I \in l_{\mathbb{R}}^2(\mathbb{Z}^2)$ such that $I[\mathbf{0}] = 1$ and $I[\mathbf{n}] = 0$ otherwise. First, for any $U, V \in l_{\mathbb{R}}^2(\mathbb{Z}^2)$ and any $s, t \in \mathbb{N}^*$, we have

$$((U \downarrow s) * V) \downarrow t = (U * (V \uparrow s)) \downarrow (st). \quad (251)$$

Then, a simple reasoning by induction yields the result, with

$$V_0^{(0)} := I; \quad V_{4l+k}^{(j+1)} := V_l^{(j)} * (G_k \uparrow 2^j) \quad (252)$$

for any $l \in \{0, \dots, j-1\}$ and any $k \in \{0, \dots, 3\}$. \square

Based on Lemma 7, the following proposition introduces complex kernels characterizing DT-CWPT.

Proposition 10. *For any $l \in \{0, \dots, 4^J - 1\}$, there exists $W_l^{\nearrow(J)} \in l_{\mathbb{C}}^2(\mathbb{Z}^2)$ such that (132) is satisfied. Identical results are obtained with the three other Fourier quadrants.*

Proof. For each of the four filter banks $m \in \{0, \dots, 3\}$, and any channel $l \in \{0, \dots, 4^J - 1\}$, Lemma 7 provides a convolution kernel $V_l^{[m](J)} \in l_{\mathbb{R}}^2(\mathbb{Z}^2)$ such that

$$D_l^{[m](J)} = (X * \bar{V}_l^{[m](J)}) \downarrow 2^J. \quad (253)$$

Then, the result is obtained by plugging (253) into (248) for all $m \in \{0, \dots, 3\}$, and by denoting

$$\begin{pmatrix} W_l^{\nearrow(J)} \\ W_l^{\searrow(J)} \end{pmatrix} = \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} V_l^{[0](J)} \\ V_l^{[3](J)} \end{pmatrix} + i \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} V_l^{[2](J)} \\ V_l^{[1](J)} \end{pmatrix}. \quad (254)$$

\square

Remark 11. DT-CWPT, computed on a discrete image X , approximates the decomposition of a continuous 2D signal $F \in L_{\mathbb{R}}^2(\mathbb{R}^2)$ into a tight frame

$$\Psi_{\mathbb{C}}^{(J)} := \bigoplus_{l=0}^{4^J-1} (\Psi_{l,\mathbf{n}}^{\nearrow(J)}, \Psi_{l,\mathbf{n}}^{\searrow(J)}, \Psi_{l,\mathbf{n}}^{\swarrow(J)}, \Psi_{l,\mathbf{n}}^{\nwarrow(J)})_{\mathbf{n} \in \mathbb{Z}^2}. \quad (255)$$

In this context, the feature maps of dual-tree wavelet packet coefficients satisfy

$$D_l^{\nearrow(J)}[\mathbf{n}] \approx (F * \bar{\Psi}_l^{\nearrow(J)*})(2^J \mathbf{n}), \quad \text{with} \quad \Psi_l^{\nearrow(J)} := \Psi_{l,\mathbf{0}}^{\nearrow(J)}. \quad (256)$$

Expression (256) is only an approximation because of implementation technicalities that occur in practice. A “perfect” dual-tree transform should be initialized with four different inputs $X^{[0-3]}$. Instead, all four WPT decompositions operate on the same input image X , leading to non-analytic outputs for small values of J . In order to counterbalance this shortcoming, the first stage of DT-CWPT decomposition must be performed with a special set of filters that satisfy the *one-sample delay* condition. We refer to Selesnick et al. (2005) for more details on this matter.

C.3 Gabor-Like Convolution Kernels

In this section, we show that the convolution kernels $W_l^{\nearrow(J)}$ and $W_l^{\searrow(J)}$, introduced in (132), approximately behave as Gabor-like filters, as defined in (11). To begin with, we assume that $h^{[0]}$ is a Shannon filter, which is associated with a sinc scaling function (Shannon, 1949). Let $J \in \mathbb{N} \setminus \{0\}$ denote the number of decomposition stages. The following proposition states that DT-CWPT tiles the frequency plane with square windows.

Proposition 11. *There exists a permutation $(\sigma_l^{(J)})_{l \in \{0, \dots, 4^J - 1\}}$ of $\{0, \dots, 2^J - 1\}^2$ such that, for any $l \in \{0, \dots, 4^J - 1\}$,*

$$\Psi_l^{\nearrow(J)} \in \mathcal{V}(\theta_l^{(J)}, \kappa_J), \quad (257)$$

where $\Psi_l^{\nearrow(J)}$ has been introduced in Remark 11, and where we have defined

$$\theta_l^{(J)} := \left(\sigma_l^{(J)} + \frac{1}{2} \right) \frac{\pi}{2^J} \quad \text{and} \quad \kappa_J := \frac{\pi}{2^J}. \quad (258)$$

We remind the reader that $\mathcal{V}(\nu, \varepsilon)$, defined in (8), denotes a space of Gabor-like filters in the continuous framework.

Proof. The atoms $\Psi_l^{\nearrow(J)}$ of the wavelet packet tight frame $\Psi_{\mathbb{C}}^{(J)}$ can be written as the tensor product of two 1D wavelet packets:

$$\Psi_l^{\nearrow(J)} = \psi_{l_1}^{(J)} \otimes \psi_{l_2}^{(J)}, \quad (259)$$

for some indices l_1 and $l_2 \in \{0, \dots, 2^J - 1\}$. Moreover, for any $l' \in \{0, \dots, 2^J - 1\}$, we have

$$\psi_{l'}^{(J)} = \psi_{l'}^{[0](J)} + i \psi_{l'}^{[1](J)}, \quad (260)$$

where $\psi_{l'}^{[0](J)} \in L_{\mathbb{R}}^2(\mathbb{R})$ is an atom of the standard Shannon wavelet packet orthonormal basis, and $\psi_{l'}^{[1](J)}$ is the one-dimensional Hilbert transform of $\psi_{l'}^{[0](J)}$. Therefore, since the Hilbert transform suppresses negative frequencies, we get

$$\widehat{\psi}_{l'}^{(J)} = 2 \widehat{\psi}_{l'}^{[0](J)} \mathbb{1}_{\mathbb{R}_+}. \quad (261)$$

Consequently, according to the Coifman-Wickerhauser theorem (Mallat, 2009, pp. 384-385), there exists $k \in \{0, \dots, 2^J - 1\}$ such that

$$\text{supp } \widehat{\psi}_{l'}^{(J)} \subset \left[\frac{k\pi}{2^J}, \frac{(k+1)\pi}{2^J} \right]. \quad (262)$$

Finally, the tensor product (259) yields the result. \square

According to Proposition 11, each atom $\Psi_l^{\nearrow(J)}$, for $l \in \{0, \dots, 4^J - 1\}$, is supported in a square window of size $\kappa_J \times \kappa_J$ included in the top-right quadrant of the Fourier domain. Similar results can be obtained for the three remaining quadrants, with $\Psi_l^{\searrow(J)}$, $\Psi_l^{\swarrow(J)}$ and $\Psi_l^{\nwarrow(J)}$. We would like to deduce from Proposition 11 that the discrete filter $W_l^{\nearrow(J)} \in l_{\mathbb{C}}^2(\mathbb{Z}^2)$ satisfies the Gabor property (133). However, as mentioned in Remark 11, (256) is only an approximation. In fact, the Fourier support of $W_l^{\nearrow(J)}$ is contained in four square regions of size κ_J (one in each quadrant), its energy becoming negligible outside the top-right quadrant when J increases. Nevertheless, employing, in the first stage, a specific pair of low-pass filters satisfying the one-sample delay condition (Selesnick et al., 2005) yields near-analytic solutions even for small values of J . We therefore assume that (133) is a reasonable approximation if $J \geq 2$.

Remark 12. Proposition 11 tiles the top-right Fourier quadrant with 4^J square cells of size $\kappa_J := \pi/2^J$. However, the Shannon wavelet is poorly suited for sparse image representations, because of its slow decay rate. Moreover, it deviates from what is typically observed in freely-trained CNNs, because $W_l^{\nearrow(J)}$ must be approximated with very large filters to avoid numerical instabilities. Practical implementations of DT-CWPT use fast-decaying filters such as these associated to Meyer wavelets (Meyer, 1985), or finite-length filters that approximate the half-sample delay condition (Selesnick et al., 2005). Therefore, energy is leaking

DEPTH J	BANDWIDTH κ_J	MEAN	STD
2	$\pi/2$	0.98	0.00
3	$\pi/4$	0.95	0.02

Table 1. Energy concentration of the DT-CWPT filters within a Fourier window of size $\kappa_J \times \kappa_J$, with $\kappa_J := \pi/2^{J-1}$.

outside the square cells tiling the Fourier domain. To counterbalance this, we increase the window size up to

$$\kappa_J := \frac{\pi}{2^{J-1}} = \pi/m_J, \quad (263)$$

and suggest that (133) remains a reasonable approximation. Therefore, the conditions to apply Theorems 1 to 3 are approximately satisfied in this context.

In order to numerically assess this assumption, we measured the maximum percentage of energy within a square window of size $\kappa_J \times \kappa_J$ in the Fourier domain:

$$\rho_i^{\nearrow} := \frac{\max_{\boldsymbol{\theta} \in [-\pi, \pi]^2} \left\| \mathbb{1}_{B_\infty(\boldsymbol{\theta}, \kappa_J/2)} \widehat{\mathbf{W}}_i^{\nearrow(J)} \right\|_{L^2}^2}{\left\| \widehat{\mathbf{W}}_i^{\nearrow(J)} \right\|_{L^2}^2}, \quad (264)$$

where the l^∞ -ball $B_\infty(\boldsymbol{\theta}, \kappa_J/2)$ is defined in the quotient space $[-\pi, \pi]^2 / (2\pi\mathbb{Z}^2)$, as explained in Remark 2. If (133) is perfectly satisfied, then $\rho_i^{\nearrow} = 1$. The statistics computed over the collection $(\rho_i^{\nearrow}, \rho_i^{\searrow})_{i \in \{0, \dots, 4^J - 1\}}$ are reported in table 1.

Remark 13. For “boundary filters”, *i.e.*, when $\|\boldsymbol{\theta}_i^{(J)}\|_\infty = (1 - 2^{-(J+1)})\pi$, Remark 2 states that a small fraction of the filter’s energy remains located at the far end of the Fourier domain—see also Bayram & Selesnick (2008). Therefore, these filters do not strictly comply with the conditions of Theorems 1 to 3. We nevertheless include them in our experiments.