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

Hubert Leterme, Kévin Polisano, Valérie Perrier, Karteek Alahari

► To cite this version:

Hubert Leterme, Kévin Polisano, Valérie Perrier, Karteek Alahari. On the Shift Invariance of Max Pooling Feature Maps in Convolutional Neural Networks. 2023. hal-03779434v2

HAL Id: hal-03779434

https://hal.science/hal-03779434v2

Preprint submitted on 24 Oct 2023

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

Hubert Leterme[†], Kévin Polisano[‡], Valérie Perrier[‡], and Karteek Alahari[†]

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.

Key words. deep learning, image processing, shift invariance, max pooling, dual-tree complex wavelet packet transform, aliasing

MSC codes. 42C40, 68U10, 94A08

1. Introduction.

Understanding the mathematical properties of deep convolutional neural networks (CNNs) [22] remains a challenging issue today. On the other hand, wavelet and multiresolution analysis are built upon a well-established mathematical framework. They have proven to be efficient for tasks such as signal compression and denoising [48], and have been widely used as feature extractors for signal, image and texture classification [17, 21, 37, 52]. There is a broad literature revealing strong connections between these two paradigms, as discussed in subsections 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 rely on convolutions and nonlinear pooling operations to transform input images into high-level feature vectors, which are in turn processed for the task at hand. In the context of image classification, the feature vectors are fed into a linear classifier. In order to achieve high classification accuracy, a convolutional network is expected to retain discriminative image components while reducing intra-class variability [9, 23]. A key property that is often desired in CNNs is their ability to remain invariant to small

---

[*] This work has been partially supported by the LabEx PERSYVAL-Lab (ANR-11-LABX-0025-01) funded by the French program Investissement d'avenir, as well as the ANR grant MIAI (ANR-19-P3IA-0003). Most of the computations presented in this paper were performed using the GRICAD infrastructure (https://gricad.univ-grenoble-alpes.fr), which is supported by Grenoble research communities.

[†] Univ. Grenoble Alpes, CNRS, Inria, Grenoble INP, LJK, 38000 Grenoble, France (hubert.leterme@univ-grenoble-alpes.fr, karteek.alahari@inria.fr).

[‡] Univ. Grenoble Alpes, CNRS, Grenoble INP, LJK, 38000 Grenoble, France (kevin.polisano@univ-grenoble-alpes.fr, valerie.perrier@univ-grenoble-alpes.fr).

input transformations, such as translations, rotations, distortions, or scaling [6, 9, 26, 43, 51]. Since perfect invariance is seldom achieved, we shall also use the term *stability* to refer to this behavior. This paper targets translations, also called shifts.

Furthermore, we focus on a configuration that is commonly observed in CNNs when trained on image datasets: many convolution kernels in the first layer resemble band-pass oriented waveforms [38, 53], referred to as *Gabor-like filters*. Whether extracted features are stable to translations is partly addressed by [2, 56]. These papers point out that strided convolution and pooling operators may greatly diverge from shift invariance, due to aliasing when subsampling high-frequency signals. In response, recent works [56, 59] introduced an antialiasing method based on low-pass filtering. They managed to increase both stability and predictive power of CNNs, despite the resulting loss of information.

In the current paper, we show that, under specific conditions that we establish, the max pooling operator can actually 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 led us to develop a method for improving shift invariance in CNNs which, unlike the previously-mentioned papers, preserves high-frequency information [25].

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 and Mallat [9] 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 [28, 29]. A variation has been proposed by Sifre and Mallat [43] 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 et al. [32, 33] 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 [44], learnable and parametric ScatterNets [10, 14], geometric ScatterNets operating on Riemanian manifolds [36], and graph ScatterNets [13, 58]. Also worth mentioning, Czaja and Li [11, 12] studied ScatterNets based on uniform covering frames, i.e., frames splitting the frequency domain into windows of roughly equal size, much like DT-$\mathbb{C}$WPT frames (as used in the present paper). Other works by Zarka et al. [54, 55] 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. [47] 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 and Bölcskei [51] 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. 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 [2, 56]. 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 [5, 7, 8, 40]. These authors showed that certain classes of CNNs are contained into 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 [3, 35, 45, 49, 57]. 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. [35].

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, $l^2_{\mathbb{R}}(\mathbb{Z}^2)$ and $l^2_{\mathbb{C}}(\mathbb{Z}^2)$ represent the discrete spaces of square-summable two-dimensional sequences with values in $\mathbb{R}$ and $\mathbb{C}$, respectively. Let

$\mathrm{W} \in l^2_{\mathbb{C}}(\mathbb{Z}^2)$ denote a two-dimensional band-pass, oriented and analytic *Gabor-like* filter, for which a formal definition will be provided in (2.5). We first consider an operator, referred to as *real-max-pooling* ($\mathbb{R}$Max), which computes the subsampled cross-correlation between an input image $\mathrm{X} \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$ and the real part of W; then calculates the maximum value over a sliding discrete grid:

$$U^{\max}_{m,\,q}[\mathrm{W}] : \mathrm{X} \mapsto \mathrm{MaxPool}_q \left(\left(\mathrm{X} * \overline{\mathrm{Re}\,\mathrm{W}}\right) \downarrow m \right), \tag{1.1}$$

where $m \in \mathbb{N} \setminus \{0\}$ denotes a subsampling factor, $\overline{\mathrm{V}}$ denotes the "flipped" sequence for any given $\mathrm{V} \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$ or $l^2_{\mathbb{C}}(\mathbb{Z}^2)$, satisfying, for any $\boldsymbol{n} \in \mathbb{Z}^2$,

$$\overline{\mathrm{V}}[\boldsymbol{n}] := \mathrm{V}[-\boldsymbol{n}], \tag{1.2}$$

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

$$(\mathrm{X} * \overline{\mathrm{V}})[\boldsymbol{n}] := \sum_{\boldsymbol{p} \in \mathbb{Z}^2} \mathrm{X}[\boldsymbol{p}]\, \overline{\mathrm{V}}[\boldsymbol{n} - \boldsymbol{p}] \qquad \text{and} \qquad (\mathrm{Y} \downarrow m)[\boldsymbol{n}] := \mathrm{Y}[m\boldsymbol{n}]. \tag{1.3}$$

In the above expression, $\mathrm{MaxPool}_q$ 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 $\mathrm{Y} \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$ and any $\boldsymbol{n} \in \mathbb{Z}^2$,

$$\mathrm{MaxPool}_q(\mathrm{Y})[\boldsymbol{n}] := \max_{\|\boldsymbol{p}\|_\infty \leq q} \mathrm{Y}[2\boldsymbol{n} + \boldsymbol{p}]. \tag{1.4}$$

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

$$U^{\mathrm{mod}}_{m}[\mathrm{W}] : \mathrm{X} \mapsto \left| (\mathrm{X} * \overline{\mathrm{W}}) \downarrow (2m) \right|. \tag{1.5}$$

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

$$U^{\mathrm{mod}}_{m}[\mathrm{W}]\,(\mathrm{X}) \approx U^{\max}_{m,\,q}[\mathrm{W}]\,(\mathrm{X}). \tag{1.6}$$

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

Remark 1.1. In the above definitions, cross-correlations are computed with a subsampling factor which is twice larger for $\mathbb{C}$Mod, compared to $\mathbb{R}$Max. However, since max pooling is also computed with subsampling, both operators have the same subsampling factor of $2m$.

We assess our theoretical findings on a deterministic setting based on the dual-tree complex wavelet packet transform (DT-$\mathbb{C}$WPT), a particular case of discrete Gabor-like decomposition with perfect reconstruction properties [4]. DT-$\mathbb{C}$WPT 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 [39]. More specifically, given an input image, we compute the mean square error between the outputs of ℂMod and ℝMax, for each wavelet packet filter. We then observe that shift invariance, when measured on ℝMax feature maps, is nearly achieved when they remain close to ℂMod outputs. We therefore establish a domain of validity for shift invariance of the ℝMax operator.

Prior to this work, we presented a preliminary study [24], where we experimentally showed that an operator based on the real part of DT-ℂWPT can mimic the behavior of the first convolution layer with fewer parameters, while keeping the network's predictive power. Our model was solely based on real-valued filters, which are prone to aliasing. Yet, it produced relatively stable outputs when compared with other models based on the standard, poorly-oriented wavelet packet transform. At the same time, we became aware of a preliminary work by Waldspurger [50, pp. 190–191], suggesting a potential connection between the combinations "real wavelet transform → max pooling" on the one hand and "complex wavelet transform → modulus" on the other hand. Following this idea, we decided to study whether invariance properties of complex moduli could somehow be captured by the max pooling operator. As shown in the present paper, Waldspurger's work does not fully extend to the discrete framework. We address this issue by adopting a probabilistic 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 $p \in \mathbb{R}_{>0} \cup \{\infty\}$, $\boldsymbol{x} \in \mathbb{R}^2$ and $r \in \mathbb{R}_+$, we denote by $B_p(\boldsymbol{x}, r) \subset \mathbb{R}^2$ the closed l^p-ball with center $\boldsymbol{x}$ and radius r. When $\boldsymbol{x} = \boldsymbol{0}$, we write $B_p(r)$.

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

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

$$\forall \boldsymbol{\nu} \in \mathbb{R}^2,\ \widehat{F}(\boldsymbol{\nu}) := \iint_{\mathbb{R}^2} F(\boldsymbol{x}) \mathrm{e}^{-i\langle \boldsymbol{\nu},\, \boldsymbol{x}\rangle}\, \mathrm{d}^2\boldsymbol{x}. \tag{2.1}$$

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

$$\mathcal{V}\big(\boldsymbol{\nu}, \varepsilon\big) := \Big\{ \Psi \in L^2_{\mathbb{C}}(\mathbb{R}^2) \;\Big|\; \operatorname{supp} \widehat{\Psi} \subset B_\infty(\boldsymbol{\nu}, \varepsilon/2) \Big\}. \tag{2.2}$$

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

$$\mathcal{T}_{\boldsymbol{h}} F : \boldsymbol{x} \mapsto F(\boldsymbol{x} - \boldsymbol{h}). \tag{2.3}$$

Discrete Framework. We denote by $l^2_{\mathbb{C}}(\mathbb{Z}^2)$ the space of 2D complex-valued square-summable sequences, represented by straight capital letters. Indexing is made between square brackets: $\forall \mathrm{X} \in l^2_{\mathbb{C}}(\mathbb{Z}^2)$, $\forall \boldsymbol{n} \in \mathbb{Z}^2$, $\mathrm{X}[\boldsymbol{n}] \in \mathbb{C}$, and we denote by $l^2_{\mathbb{R}}(\mathbb{Z}^2) \subset l^2_{\mathbb{C}}(\mathbb{Z}^2)$ the subset of real-valued sequences. For any $\mathrm{V} \in l^2_{\mathbb{C}}(\mathbb{Z}^2)$, $\overline{\mathrm{V}}$ denotes its "flipped" version as defined in (1.2). The convolution and subsampling operators, respectively denoted by $*$ and $\downarrow$, are defined in (1.3). 2D images, feature maps and convolution kernels are considered as elements of $l^2_{\mathbb{C}}(\mathbb{Z}^2)$. Besides, multichannel arrays of 2D sequences are denoted by bold straight capital letters, for instance: $\mathbf{X} := (\mathrm{X}_k)_{k\in\{0..K-1\}}$. Note that indexing starts at 0 to comply with practical implementations.

The 2D discrete-time Fourier transform of any $\mathrm{X} \in l^2_{\mathbb{C}}(\mathbb{Z}^2)$, denoted by $\widehat{\mathrm{X}} \in L^2_{\mathbb{C}}([-\pi,\, \pi]^2)$, is defined by

$$\forall \boldsymbol{\theta} \in [-\pi,\, \pi]^2\,,\ \widehat{\mathrm{X}}(\boldsymbol{\theta}) := \sum_{\boldsymbol{n}\in\mathbb{Z}^2} \mathrm{X}[\boldsymbol{n}]\mathrm{e}^{-i\langle\boldsymbol{\theta},\,\boldsymbol{n}\rangle}. \tag{2.4}$$

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

$$\mathcal{J}\big(\boldsymbol{\theta},\, \kappa\big) := \Big\{\mathrm{W} \in l^2_{\mathbb{C}}(\mathbb{Z}^2) \;\Big|\; \operatorname{supp}\widehat{\mathrm{W}} \subset B_\infty(\boldsymbol{\theta},\, \kappa/2)\Big\}\,. \tag{2.5}$$

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

Remark 2.1. 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.

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 $\mathrm{V} \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$ one of these "well-behaved" filters. Its Fourier spectrum roughly consists in two bright spots which are symmetric with respect to the origin.[1] Now, we consider a complex-valued companion $\mathrm{W} \in l^2_{\mathbb{C}}(\mathbb{Z}^2)$ such that

$$\widehat{\mathrm{W}}(\boldsymbol{\omega}) := \big(1 + \operatorname{sgn}\langle\boldsymbol{\omega},\, \boldsymbol{u}\rangle\big) \cdot \widehat{\mathrm{V}}(\boldsymbol{\omega}) \qquad \forall\boldsymbol{\omega} \in [-\pi,\, \pi]^2\,, \tag{2.6}$$

where $\boldsymbol{u}$ denotes a unit vector orthogonal to the filter's orientation.

We can show that V is the real part of W, and that $\mathrm{W} = \mathrm{V} + i\mathcal{H}(\mathrm{V})$, where $\mathcal{H}$ denotes the two-dimensional Hilbert transform as introduced by Havlicek et al. [16]. It satisfies

$$\widehat{\mathcal{H}(\mathrm{V})}(\boldsymbol{\omega}) := -i\operatorname{sgn}\langle\boldsymbol{\omega},\, \boldsymbol{u}\rangle\, \widehat{\mathrm{V}}(\boldsymbol{\omega}). \tag{2.7}$$

---

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

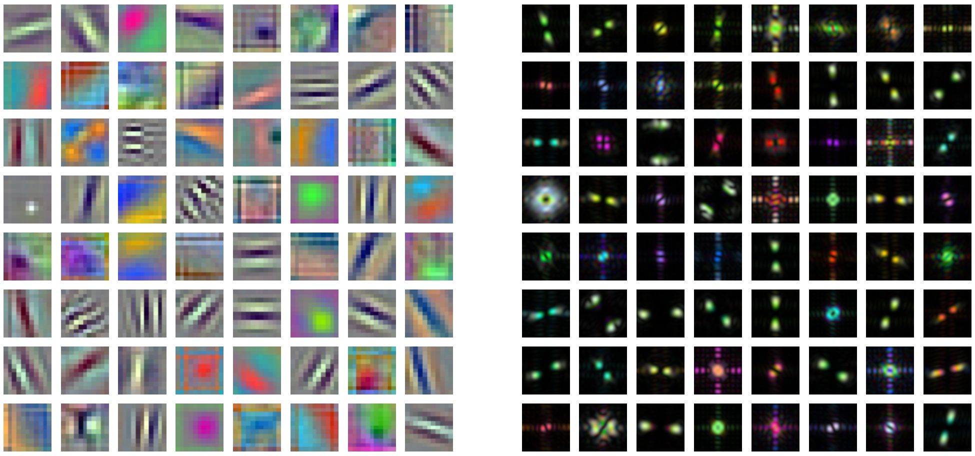

Figure 1. Spatial (left) and Fourier (right) representations of convolution kernels in the first layer of AlexNet, after training with ImageNet ILSVRC 2012-2017 [39]. Each kernel connects the 3 RGB input channels to one of the 64 output channels.

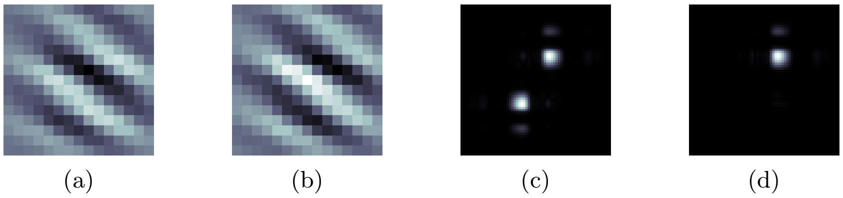

(a) (b) (c) (d)

Figure 2. (a), (b): Real and imaginary parts of a Gabor-like filter W as defined in (2.6). (c), (d): Magnitude spectra (modulus of the Fourier transform) of V and W, respectively.

As a consequence, $\widehat{\mathrm{W}}$ is equal to $2\widehat{\mathrm{V}}$ 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 Figure 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.

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 $\mathrm{W} \in l^2_{\mathbb{C}}(\mathbb{Z}^2)$, $\mathbb{C}$Mod is nearly shift-invariant. To clarify, we establish that

$$U_m^{\mathrm{mod}}[\mathrm{W}]\,(\mathrm{X}) \approx U_m^{\mathrm{mod}}[\mathrm{W}]\,(\mathcal{T}_{\boldsymbol{u}}\mathrm{X}), \tag{2.8}$$

for “small” translation vectors $\boldsymbol{u} \in \mathbb{R}^2$, where a formal definition of the translation operator will be defined in (2.34). This result is hinted by Kingsbury and Magarey [20] 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 2.2 below is adapted from Waldspurger [50, pp. 190–191].

Lemma 2.2. *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* (2.2). *Now, for any real-valued function* $F \in L^2_{\mathbb{R}}(\mathbb{R}^2)$, *we consider the complex-valued function* $F_0 \in L^2_{\mathbb{C}}(\mathbb{R}^2)$ *defined by*

$$F_0 : \boldsymbol{x} \mapsto (F * \overline{\Psi})(\boldsymbol{x})\, \mathrm{e}^{i\langle \boldsymbol{\nu},\, \boldsymbol{x}\rangle}. \tag{2.9}$$

Then F_0 *is low-frequency. Specifically,*

$$\operatorname{supp} \widehat{F_0} \subset B_\infty(\varepsilon/2). \tag{2.10}$$

Proof. Applying the Fourier transform on (2.9) yields, for any $\boldsymbol{\xi} \in \mathbb{R}^2$,

$$\widehat{F_0}(\boldsymbol{\xi}) = \big(\widehat{F * \overline{\Psi}}\big)(\boldsymbol{\xi} - \boldsymbol{\nu}) = \mathcal{T}_{\boldsymbol{\nu}}\big(\widehat{F}\,\widehat{\overline{\Psi}}\big)(\boldsymbol{\xi}). \tag{2.11}$$

By hypothesis on Ψ, we have

$$\operatorname{supp}(\widehat{F}\,\widehat{\overline{\Psi}}) \subset \operatorname{supp} \widehat{\overline{\Psi}} \subset B_\infty(-\boldsymbol{\nu},\, \varepsilon/2). \tag{2.12}$$

The translation operator $\mathcal{T}_{\boldsymbol{\nu}}$ shifts the support with respect to $\boldsymbol{\nu}$, which yields (2.10). ■

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

Proposition 2.3. *For any* $F_0 \in L^2_{\mathbb{R}}(\mathbb{R}^2)$ *such that* $\operatorname{supp} \widehat{F_0} \subset B_\infty(\varepsilon/2)$, *and any* $\boldsymbol{h} \in \mathbb{R}^2$,

$$\|\mathcal{T}_{\boldsymbol{h}} F_0 - F_0\|_{L^2} \leq \alpha(\varepsilon \boldsymbol{h})\, \|F_0\|_{L^2}\,, \tag{2.13}$$

where we have defined

$$\alpha : \boldsymbol{\tau} \mapsto \frac{\|\boldsymbol{\tau}\|_1}{2}. \tag{2.14}$$

Proof. Using the 2D Plancherel formula, we compute

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

because $\cos t \geq 1 - \frac{t^2}{2}$. Note that the integral is computed on a compact domain because, according to Lemma 2.2, $\operatorname{supp} \widehat{F_0} \subset B_\infty(\varepsilon/2)$. Now, we use the Cauchy-Schwarz inequality to compute:

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

Therefore,

$$\|\mathcal{T}_{\boldsymbol{h}} F_0 - F_0\|_{L^2}^2 \leq \frac{\varepsilon}{4} \|\boldsymbol{h}\|_1^2 \, \|F_0\|_{L^2}^2 , \tag{2.15}$$

which yields the result. ■

2.4. Adaptation to Discrete 2D Sequences. Given $\kappa \in \,]0,\, 2\pi]$ and $\boldsymbol{\theta} \in B_\infty(\pi)$, let $\mathrm{W} \in \mathcal{J}\big(\boldsymbol{\theta},\, \kappa\big)$ denote a discrete Gabor-like filter such as defined in (2.5). For any image $\mathrm{X} \in l^2_{\mathbb{C}}(\mathbb{Z}^2)$ with finite support and any subsampling factor $m \in \mathbb{N}\backslash\{0\}$, we express $(\mathrm{X} * \overline{\mathrm{W}}) \downarrow m$ using the continuous framework introduced above, and derive an invariance formula.

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

$$\widehat{\varPhi^{(s)}} := s \mathbb{1}_{B_\infty(\pi/s)}. \tag{2.16}$$

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

$$\varPhi^{(s)}_{\boldsymbol{n}}(\boldsymbol{x}) := \varPhi^{(s)}(\boldsymbol{x} - s\boldsymbol{n}). \tag{2.17}$$

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

$$\mathcal{V}^{(s)} := \big\{F \in L^2_{\mathbb{C}}(\mathbb{R}^2) \,\big|\, \operatorname{supp} \widehat{F} \subset B_\infty(\pi/s)\big\}. \tag{2.18}$$

Then, using the notation introduced in (2.2), we have $\mathcal{V}^{(s)} = \mathcal{V}(\mathbf{0},\, 2\pi/s)$.

We now consider the following lemma.

Lemma 2.4. *Let $s > 0$. For any $F \in \mathcal{V}^{(s)}$ and any $\boldsymbol{\xi} \in B_\infty(\pi/s)$, we have*

$$\widehat{F}(\boldsymbol{\xi}) = s\, \widehat{\mathrm{X}}(s\boldsymbol{\xi}), \tag{2.19}$$

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

$$\|F\|_{L^2} = \|\mathrm{X}\|_2 . \tag{2.20}$$

Proof. Since $F \in \mathcal{V}^{(s)}$, the two-dimensional version of Shannon's sampling theorem [27, Theorem 3.11, p. 81] yields

$$F = \sum_{\boldsymbol{n}\in\mathbb{Z}^2} \mathrm{X}[\boldsymbol{n}]\, \Phi_{\boldsymbol{n}}^{(s)}, \qquad \text{and} \qquad \widehat{F} = \sum_{\boldsymbol{n}\in\mathbb{Z}^2} \mathrm{X}[\boldsymbol{n}]\, \widehat{\Phi_{\boldsymbol{n}}^{(s)}}. \tag{2.21}$$

Besides, using (2.16), we can show that, for any $\boldsymbol{\xi} \in B_\infty(\pi/s)$,

$$\widehat{\Phi_{\boldsymbol{n}}^{(s)}}(\boldsymbol{\xi}) = \widehat{\Phi^{(s)}}(\boldsymbol{\xi})\, \mathrm{e}^{-i\langle s\boldsymbol{\xi},\, \boldsymbol{n}\rangle} = s\, \mathrm{e}^{-i\langle s\boldsymbol{\xi},\, \boldsymbol{n}\rangle}. \tag{2.22}$$

Therefore, plugging (2.22) into (2.21) proves (2.19).

Then, by combining (2.19) with the Plancherel formula, we get

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

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

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

hence (2.20), which concludes the proof. ■

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

Proposition 2.5. *Let* $\mathrm{X} \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$ *denote an input image with finite support, and* $\mathrm{W} \in \mathcal{J}\big(\boldsymbol{\theta},\, \kappa\big)$. *Considering a sampling interval* $s \in \mathbb{R}_{>0}$, *we define* F_{X} *and* $\Psi_{\mathrm{W}} \in \mathcal{V}^{(s)}$ *such that*

$$F_{\mathrm{X}} := \sum_{\boldsymbol{n}\in\mathbb{Z}^2} \mathrm{X}[\boldsymbol{n}]\, \Phi_{\boldsymbol{n}}^{(s)} \qquad \textit{and} \qquad \Psi_{\mathrm{W}} := \sum_{\boldsymbol{n}\in\mathbb{Z}^2} \mathrm{W}[\boldsymbol{n}]\, \Phi_{\boldsymbol{n}}^{(s)}. \tag{2.23}$$

Then,

$$\Psi_{\mathrm{W}} \in \mathcal{V}\big(\boldsymbol{\theta}/s,\, \kappa/s\big). \tag{2.24}$$

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

$$\mathrm{X}[\boldsymbol{n}] = s\, F_{\mathrm{X}}(s\boldsymbol{n}); \qquad \mathrm{W}[\boldsymbol{n}] = s\, \Psi_{\mathrm{W}}(s\boldsymbol{n}), \tag{2.25}$$

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

$$\big((\mathrm{X} * \overline{\mathrm{W}}) \downarrow m\big)\, [\boldsymbol{n}] = \big(F_{\mathrm{X}} * \overline{\Psi}_{\mathrm{W}}\big)\, (ms\boldsymbol{n})\,. \tag{2.26}$$

Proof. First, F_{X} and $\mathit{\Psi}_{\mathrm{W}}$ are well defined because $\mathrm{X} \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$ and $\mathrm{W} \in l^2_{\mathbb{C}}(\mathbb{Z}^2)$. By construction, F_{X} and $\mathit{\Psi}_{\mathrm{W}} \in \mathcal{V}^{(s)}$. Therefore, according to Shannon's sampling theorem [27, Theorem 3.11, p. 81],

$$F_{\mathrm{X}} := s \sum_{\boldsymbol{n}\in\mathbb{Z}^2} F_{\mathrm{X}}(s\boldsymbol{n})\, \mathit{\Phi}^{(s)}_{\boldsymbol{n}} \qquad \text{and} \qquad \mathit{\Psi}_{\mathrm{W}} := s \sum_{\boldsymbol{n}\in\mathbb{Z}^2} \mathit{\Psi}_{\mathrm{W}}(s\boldsymbol{n})\, \mathit{\Phi}^{(s)}_{\boldsymbol{n}}. \tag{2.27}$$

By uniqueness of decompositions in an orthonormal basis, we get (2.25). Moreover, using (2.19) in Lemma 2.4, we get, for any $\boldsymbol{\xi} \in B_\infty(\pi/s)$,

$$\widehat{\mathit{\Psi}}_{\mathrm{W}}(\boldsymbol{\xi}) = s\, \widehat{\mathrm{W}}(s\boldsymbol{\xi}). \tag{2.28}$$

Since $\widehat{\mathit{\Psi}}_{\mathrm{W}}(\boldsymbol{\xi}) = 0$ outside $B_\infty(\pi/s)$, (2.28) is true for any $\boldsymbol{\xi} \in \mathbb{R}^2$. Therefore, by hypothesis on W,

$$\operatorname{supp} \widehat{\mathit{\Psi}}_{\mathrm{W}} \subset B_\infty\big(\boldsymbol{\theta}/s,\, \kappa/(2s)\big), \tag{2.29}$$

which yields (2.24).

We now prove (2.26). For $\boldsymbol{n} \in \mathbb{Z}^2$, we compute:

$$\begin{aligned}
\left(F_{\mathrm{X}} * \overline{\mathit{\Psi}}_{\mathrm{W}}\right)(ms\boldsymbol{n}) &= \iint_{\mathbb{R}^2} F_{\mathrm{X}}(ms\boldsymbol{n} - \boldsymbol{x})\, \overline{\mathit{\Psi}}_{\mathrm{W}}(\boldsymbol{x})\, \mathrm{d}^2\boldsymbol{x} \\
&= \iint_{\mathbb{R}^2} \sum_{\boldsymbol{p}\in\mathbb{Z}^2} \mathrm{X}[\boldsymbol{p}] \mathit{\Phi}^{(s)}_{\boldsymbol{p}}(ms\boldsymbol{n} - \boldsymbol{x})\, \overline{\mathit{\Psi}}_{\mathrm{W}}(\boldsymbol{x})\, \mathrm{d}^2\boldsymbol{x} \\
&= \sum_{\boldsymbol{p}\in\mathbb{Z}^2} \mathrm{X}[\boldsymbol{p}] \iint_{\mathbb{R}^2} \mathit{\Phi}^{(s)}_{\boldsymbol{p}}(ms\boldsymbol{n} - \boldsymbol{x})\, \overline{\mathit{\Psi}}_{\mathrm{W}}(\boldsymbol{x})\, \mathrm{d}^2\boldsymbol{x}.
\end{aligned}$$

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

$$\begin{aligned}
\left(F_{\mathrm{X}} * \overline{\mathit{\Psi}}_{\mathrm{W}}\right)(ms\boldsymbol{n}) &= \sum_{\boldsymbol{p}\in\mathbb{Z}^2} \mathrm{X}[\boldsymbol{p}] \iint_{\mathbb{R}^2} \overline{\mathit{\Psi}}_{\mathrm{W}}(\boldsymbol{x})\, \mathit{\Phi}^{(s)}\big(s(m\boldsymbol{n} - \boldsymbol{p}) - \boldsymbol{x}\big)\, \mathrm{d}^2\boldsymbol{x} && (2.30) \\
&= \sum_{\boldsymbol{p}\in\mathbb{Z}^2} \mathrm{X}[\boldsymbol{p}] \left(\overline{\mathit{\Psi}}_{\mathrm{W}} * \mathit{\Phi}^{(s)}\right)\big(s(m\boldsymbol{n} - \boldsymbol{p})\big) && (2.31)
\end{aligned}$$

Since $\{\mathit{\Phi}^{(s)}_{\boldsymbol{n}}\}_{\boldsymbol{n}\in\mathbb{Z}^2}$ is an orthonormal basis of $\mathcal{V}^{(s)}$, the definition of $\mathit{\Psi}_{\mathrm{W}}$ in (2.23) implies, for any $\boldsymbol{p}' \in \mathbb{Z}^2$,

$$\overline{\mathrm{W}}[\boldsymbol{p}'] = \left\langle \mathit{\Psi}_{\mathrm{W}},\, \mathit{\Phi}^{(s)}_{-\boldsymbol{p}'} \right\rangle = \left(\overline{\mathit{\Psi}}_{\mathrm{W}} * \mathit{\Phi}^{(s)}\right)(s\boldsymbol{p}'), \tag{2.32}$$

because $\mathit{\Phi}^{(s)}$ is even. Therefore, plugging (2.32) with $\boldsymbol{p}' \leftarrow (m\boldsymbol{n} - \boldsymbol{p})$ into (2.31) yields

$$\left(F_{\mathrm{X}} * \overline{\mathit{\Psi}}_{\mathrm{W}}\right)(ms\boldsymbol{n}) = \sum_{\boldsymbol{p}\in\mathbb{Z}^2} \mathrm{X}[\boldsymbol{p}]\, \overline{\mathrm{W}}[m\boldsymbol{n} - \boldsymbol{p}] = \big(\mathrm{X} * \overline{\mathrm{W}}\big)[m\boldsymbol{n}], \tag{2.33}$$

hence the result. ■

Proposition 2.5 introduces a latent subspace of $L^2_{\mathbb{R}}(\mathbb{R}^2)$ from which input images are uniformly sampled. This allows us to define, for any $\boldsymbol{u} \in \mathbb{R}^2$, a translation operator $\mathcal{T}_{\boldsymbol{u}}$ on discrete sequences, even if $\boldsymbol{u}$ has non-integer values:

$$\mathcal{T}_{\boldsymbol{u}}\mathrm{X}[\boldsymbol{n}] := s\, \mathcal{T}_{s\boldsymbol{u}} F_{\mathrm{X}}(s\boldsymbol{n}), \tag{2.34}$$

where F_{X} is defined in (2.23). We can indeed show that this definition is independent from the choice of sampling interval $s > 0$. Besides, given $\mathrm{X} \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$, we have

$$\forall \boldsymbol{p} \in \mathbb{Z}^2,\ \mathcal{T}_{\boldsymbol{p}}\mathrm{X}[\boldsymbol{n}] = \mathrm{X}[\boldsymbol{n} - \boldsymbol{p}]; \tag{2.35}$$

$$\forall \boldsymbol{u},\, \boldsymbol{v} \in \mathbb{R}^2,\ \mathcal{T}_{\boldsymbol{u}}(\mathcal{T}_{\boldsymbol{v}}\mathrm{X}) = \mathcal{T}_{\boldsymbol{u}+\boldsymbol{v}}\mathrm{X}, \tag{2.36}$$

which shows that $\mathcal{T}_{\boldsymbol{u}}$ corresponds to the intuitive idea of a translation operator. Expressions (2.35) and (2.36) are direct consequence of the following lemma, which bonds the shift operator in the discrete and continuous frameworks.

Lemma 2.6. *For any* $\mathrm{X} \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$ *and any* $\boldsymbol{u} \in \mathbb{R}^2$*,*

$$F_{\mathcal{T}_{\boldsymbol{u}}\mathrm{X}} = \mathcal{T}_{s\boldsymbol{u}} F_{\mathrm{X}}. \tag{2.37}$$

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

$$F_{\mathcal{T}_{\boldsymbol{u}}\mathrm{X}} = s \sum_{\boldsymbol{n} \in \mathbb{Z}^2} \mathcal{T}_{s\boldsymbol{u}} F_{\mathrm{X}}(s\boldsymbol{n})\, \varPhi^{(s)}_{\boldsymbol{n}}. \tag{2.38}$$

On the other hand, $F_{\mathrm{X}} \in \mathcal{V}^{(s)}$ by construction. Therefore, $\mathcal{T}_{s\boldsymbol{u}} F_{\mathrm{X}} \in \mathcal{V}^{(s)}$. Then, according to Shannon's sampling theorem [27, Theorem 3.11, p. 81], we get

$$\mathcal{T}_{s\boldsymbol{u}} F_{\mathrm{X}} = s \sum_{\boldsymbol{n} \in \mathbb{Z}^2} \mathcal{T}_{s\boldsymbol{u}} F_{\mathrm{X}}(s\boldsymbol{n})\, \varPhi^{(s)}_{\boldsymbol{n}}, \tag{2.39}$$

which concludes the proof. ■

We now consider the following corollary to Proposition 2.5.

Corollary 2.7. *For any shift vector* $\boldsymbol{u} \in \mathbb{R}^2$*, we have*

$$\big((\mathcal{T}_{\boldsymbol{u}}\mathrm{X} * \overline{\mathrm{W}}) \downarrow m\big)\, [\boldsymbol{n}] = \big(\mathcal{T}_{s\boldsymbol{u}} F_{\mathrm{X}} * \overline{\varPsi}_{\mathrm{W}}\big)\, (ms\boldsymbol{n})\,. \tag{2.40}$$

Proof. Applying (2.26) in Proposition 2.5 with $\mathrm{X} \leftarrow \mathcal{T}_{\boldsymbol{u}}\mathrm{X}$, we get

$$\big((\mathcal{T}_{\boldsymbol{u}}\mathrm{X} * \overline{\mathrm{W}}) \downarrow m\big)\, [\boldsymbol{n}] = \big(F_{\mathcal{T}_{\boldsymbol{u}}\mathrm{X}} * \overline{\varPsi}_{\mathrm{W}}\big)\, (ms\boldsymbol{n})\,, \tag{2.41}$$

and Lemma 2.6 concludes the proof. ■

2.5. Shift Invariance in the Discrete Framework. We consider the $\mathbb{C}$Mod operator defined in (1.5). For the sake of conciseness, in what follows we will write U_m^{mod} instead of $U_m^{\mathrm{mod}}[\mathrm{W}]$, when no ambiguity is possible. First, we state the following lemma.

Lemma 2.8. *For any input image* $\mathrm{X} \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$ *with finite support, and any Gabor-like filter* $\mathrm{W} \in \mathcal{J}\big(\boldsymbol{\theta},\, \kappa\big)$, *we consider the low-frequency function*

$$F_0 : \boldsymbol{x} \mapsto (F_{\mathrm{X}} * \overline{\Psi}_{\mathrm{W}})(\boldsymbol{x})\, \mathrm{e}^{i\langle \boldsymbol{\theta}/s,\, \boldsymbol{x}\rangle}, \tag{2.42}$$

with F_{X} *and* Ψ_{W} *satisfying* (2.23). *If* $\kappa \leq \pi/m$, *then*

$$F_0 \in \mathcal{V}^{(s')}. \tag{2.43}$$

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

$$\sum_{\boldsymbol{n}\in\mathbb{Z}^2} \Big|\mathcal{T}_{\boldsymbol{h}} F_0(s'\boldsymbol{n}) - F_0(s'\boldsymbol{n})\Big|^2 = \frac{1}{s'^2}\, \|\mathcal{T}_{\boldsymbol{h}} F_0 - F_0\|^2_{L^2}\,, \tag{2.44}$$

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

$$\big\|U_m^{\mathrm{mod}}\mathrm{X}\big\|_2 = \frac{1}{s'}\, \|F_0\|_{L^2}\,. \tag{2.45}$$

Proof. Let us write:

$$\sum_{\boldsymbol{n}\in\mathbb{Z}^2} \big|\mathcal{T}_{\boldsymbol{h}} F_0(s'\boldsymbol{n}) - F_0(s'\boldsymbol{n})\big|^2 = \sum_{\boldsymbol{n}\in\mathbb{Z}^2} \big|F_1(s'\boldsymbol{n})\big|^2 = \frac{1}{s'^2}\, \|\mathrm{X}_1\|_2^2\,, \tag{2.46}$$

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

$$F_1 := \mathcal{T}_{\boldsymbol{h}} F_0 - F_0 \qquad \text{and} \qquad \mathrm{X}_1[\boldsymbol{n}] := s' F_1(s'\boldsymbol{n}). \tag{2.47}$$

According to Proposition 2.5 (2.24), $\Psi_{\mathrm{W}} \in \mathcal{V}\big(\boldsymbol{\theta}/s,\, \kappa/s\big)$. Therefore, according to Lemma 2.2,

$$\operatorname{supp} \widehat{F_0} \subset B_\infty\left(\frac{\kappa}{2s}\right). \tag{2.48}$$

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

$$B_\infty\left(\frac{\kappa}{2s}\right) \subset B_\infty\left(\frac{\pi}{s'}\right), \tag{2.49}$$

which yields (2.43), and $F_1 \in \mathcal{V}^{(s')}$. Then, according to Lemma 2.4 (2.20) with $\mathrm{X} \leftarrow \mathrm{X}_1$, $F \leftarrow F_1$ and $s \leftarrow s'$,

$$\|\mathrm{X}_1\|_2 = \|F_1\|_{L^2} = \|\mathcal{T}_{\boldsymbol{h}} F_0 - F_0\|_{L^2}\,. \tag{2.50}$$

Therefore, plugging (2.50) into (2.46) yields (2.44).

Besides, according again to Lemma 2.4,

$$\|F_0\|_{L^2}^2 = \|\mathrm{X}_0\|_2^2\,, \tag{2.51}$$

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

$$\mathrm{X}_0[\boldsymbol{n}] := s' F_0(s'\boldsymbol{n}). \tag{2.52}$$

Then,

$$\begin{aligned} \|\mathrm{X}_0\|_2^2 &= s'^2 \sum_{\boldsymbol{n}\in\mathbb{Z}^2} \big|\big(F_{\mathrm{X}} * \overline{\varPsi}_{\mathrm{W}}\big)(s'\boldsymbol{n})\big|^2 && \text{(acc. to (2.42))} \\ &= s'^2 \sum_{\boldsymbol{n}\in\mathbb{Z}^2} \big|(\mathrm{X} * \overline{\mathrm{W}}) \downarrow (2m)[\boldsymbol{n}]\big|^2 && \text{(acc. to Proposition 2.5 with } m \leftarrow 2m) \\ &= s'^2 \big\|U_m^{\mathrm{mod}}\mathrm{X}\big\|_2^2. && \text{(acc. to (1.5))} \end{aligned}$$

Finally, plugging this result into (2.51) yields (2.45) and concludes the proof. ■

We are now ready to state the main result about shift invariance of ℂMod outputs.

Theorem 2.9 (Shift invariance of ℂMod). *Let* $\mathrm{W} \in \mathcal{J}\big(\boldsymbol{\theta},\, \kappa\big)$ *denote a discrete Gabor-like filter and* $m \in \mathbb{N} \setminus \{0\}$ *denote a subsampling factor. Then, under the following condition:*

$$\kappa \leq \pi/m, \tag{2.53}$$

we have, for any input image $\mathrm{X} \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$ *with finite support and any translation vector* $\boldsymbol{u} \in \mathbb{R}^2$*,*

$$\big\|U_m^{\mathrm{mod}}(\mathcal{T}_{\boldsymbol{u}}\mathrm{X}) - U_m^{\mathrm{mod}}\mathrm{X}\big\|_2 \leq \alpha(\kappa\boldsymbol{u})\, \big\|U_m^{\mathrm{mod}}\mathrm{X}\big\|_2, \tag{2.54}$$

where α *has been defined in* (2.14).

Proof. As in Lemma 2.8, we consider the low-frequency function F_0 satisfying (2.42), and denote $s' := 2ms$. We can write

$$|F_{\mathrm{X}} * \overline{\varPsi}_{\mathrm{W}}| = |F_0| \qquad \text{and} \qquad |\mathcal{T}_{s\boldsymbol{u}}F_{\mathrm{X}} * \overline{\varPsi}_{\mathrm{W}}| = |\mathcal{T}_{s\boldsymbol{u}}F_0|. \tag{2.55}$$

Recall that $U_m^{\mathrm{mod}}\mathrm{X} = \big|(\mathrm{X} * \overline{\mathrm{W}}) \downarrow (2m)\big|$, such as defined in (1.5). According to Proposition 2.5 (2.26) and Corollary 2.7 (2.40) with $m \leftarrow 2m$, we therefore get

$$U_m^{\mathrm{mod}}\mathrm{X}[\boldsymbol{n}] = \big|F_0(s'\boldsymbol{n})\big|\,; \tag{2.56}$$

$$U_m^{\mathrm{mod}}(\mathcal{T}_{\boldsymbol{u}}\mathrm{X})[\boldsymbol{n}] = \big|(\mathcal{T}_{s\boldsymbol{u}}F_0)(s'\boldsymbol{n})\big|\,. \tag{2.57}$$

Then, using (2.56), (2.57) and the reverse triangle inequality,

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

Since condition (2.53) is satisfied, we can use Lemma 2.8 (2.44) with $\boldsymbol{h} \leftarrow s\boldsymbol{u}$:

$$\left\| U_m^{\mathrm{mod}}(\mathcal{T}_{\boldsymbol{u}}\mathrm{X}) - U_m^{\mathrm{mod}}\mathrm{X} \right\|_2^2 \leq \frac{1}{s'^2} \left\| \mathcal{T}_{s\boldsymbol{u}}F_0 - F_0 \right\|_{L^2}^2 \tag{2.58}$$

Now, according to Proposition 2.3 with $\varepsilon \leftarrow \kappa/s$ and $\boldsymbol{h} \leftarrow s\boldsymbol{u}$, we then get the following bound:

$$\left\| U_m^{\mathrm{mod}}(\mathcal{T}_{\boldsymbol{u}}\mathrm{X}) - U_m^{\mathrm{mod}}\mathrm{X} \right\|_2^2 \leq \frac{\alpha(\kappa\boldsymbol{u})^2}{s'^2} \left\| F_0 \right\|_{L^2}^2 . \tag{2.59}$$

Finally, using Lemma 2.8 (2.45) yields (2.54), which completes the proof. ■

Interestingly, the reference value used in Theorem 2.9, *i.e.*, $\left\| U_m^{\mathrm{mod}}\mathrm{X} \right\|_2$, is fully shift-invariant, as stated in the following proposition.

Proposition 2.10. *Let* $\mathrm{W} \in \mathcal{J}\big(\boldsymbol{\theta}, \kappa\big)$ *and* $m \in \mathbb{N} \setminus \{0\}$*. Under condition* (2.53)*, we have, for any* $\mathrm{X} \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$ *and any* $\boldsymbol{u} \in \mathbb{R}^2$*,*

$$\left\| U_m^{\mathrm{mod}}(\mathcal{T}_{\boldsymbol{u}}\mathrm{X}) \right\|_2 = \left\| U_m^{\mathrm{mod}}\mathrm{X} \right\|_2 . \tag{2.60}$$

Proof. Let $\mathrm{X} \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$ and $s > 0$. We consider $F_0 \in L^2_{\mathbb{C}}(\mathbb{R}^2)$ as the “low-frequency” function satisfying (2.42). Again, we introduce $s' := 2ms$ and $\mathrm{X}_0 \in l^2_{\mathbb{C}}(\mathbb{Z}^2)$ satisfying (2.52). Moreover, for any $\mathrm{Y} \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$, we denote by $F^{(s')}_{\mathrm{Y}}$ the Shannon interpolation of Y parameterized by s', analogously to (2.23):

$$F^{(s')}_{\mathrm{Y}} := \sum_{\boldsymbol{n}\in\mathbb{Z}^2} \mathrm{Y}[\boldsymbol{n}]\, \Phi^{(s')}_{\boldsymbol{n}} . \tag{2.61}$$

On the one hand, Lemma 2.8 provides (2.45). On the other hand, we seek a similar result with $\mathrm{X} \leftarrow \mathcal{T}_{\boldsymbol{u}}\mathrm{X}$. For this purpose, (2.57) can be rewritten

$$U_m^{\mathrm{mod}}(\mathcal{T}_{\boldsymbol{u}}\mathrm{X})[\boldsymbol{n}] = \left| \mathcal{T}_{s'\boldsymbol{u}'}F_0(s'\boldsymbol{n}) \right|, \tag{2.62}$$

with $\boldsymbol{u}' := \boldsymbol{u}/(2m)$. Besides, according to Lemma 2.8 (2.43), $F_0 \in \mathcal{V}^{(s')}$. Therefore, Shannon's sampling theorem [27, Theorem 3.11, p. 81] with $s \leftarrow s'$ yields

$$\begin{aligned} F_0 &= s' \sum_{\boldsymbol{n}\in\mathbb{Z}^2} F_0(s'\boldsymbol{n})\, \Phi^{(s')}_{\boldsymbol{n}} \\ &= \sum_{\boldsymbol{n}\in\mathbb{Z}^2} \mathrm{X}_0[\boldsymbol{n}]\, \Phi^{(s')}_{\boldsymbol{n}} = F^{(s')}_{\mathrm{X}_0}, \end{aligned}$$

where we have used the notations introduced in (2.52) and (2.61). Then, using Lemma 2.6 with $\mathrm{X} \leftarrow \mathrm{X}_0$, $\boldsymbol{u} \leftarrow \boldsymbol{u}'$ and $s \leftarrow s'$, we get

$$F^{(s')}_{\mathcal{T}_{\boldsymbol{u}'}\mathrm{X}_0} = \mathcal{T}_{s'\boldsymbol{u}'}F^{(s')}_{\mathrm{X}_0} = \mathcal{T}_{s'\boldsymbol{u}'}F_0 . \tag{2.63}$$

Besides, (2.25) (from Proposition 2.5) with $\mathrm{X} \leftarrow \mathcal{T}_{\boldsymbol{u}'}\mathrm{X}_0$ and $s \leftarrow s'$ becomes

$$\mathcal{T}_{\boldsymbol{u}'}\mathrm{X}_0[\boldsymbol{n}] = s'\, F^{(s')}_{\mathcal{T}_{\boldsymbol{u}'}\mathrm{X}_0}(s'\boldsymbol{n}), \tag{2.64}$$

and inserting (2.63) into (2.64) yields

$$\mathcal{T}_{\boldsymbol{u}'}\mathrm{X}_0[\boldsymbol{n}] = s'\,\mathcal{T}_{s'\boldsymbol{u}'}F_0(s'\boldsymbol{n}). \tag{2.65}$$

Therefore, (2.62) and (2.65) imply

$$\left\|U_m^{\mathrm{mod}}(\mathcal{T}_{\boldsymbol{u}}\mathrm{X})\right\|_2 = \frac{1}{s'}\left\|\mathcal{T}_{\boldsymbol{u}'}\mathrm{X}_0\right\|_2. \tag{2.66}$$

Moreover, since $F_0 \in \mathcal{V}^{(s')}$, and according to (2.65), we can use Lemma 2.4 with $s \leftarrow s'$, $F \leftarrow \mathcal{T}_{s'\boldsymbol{u}'}F_0$ and $\mathrm{X} \leftarrow \mathcal{T}_{\boldsymbol{u}'}\mathrm{X}_0$. We get

$$\left\|\mathcal{T}_{\boldsymbol{u}'}\mathrm{X}_0\right\|_2 = \left\|\mathcal{T}_{s'\boldsymbol{u}'}F_0\right\|_{L^2} = \left\|F_0\right\|_{L^2}, \tag{2.67}$$

and plugging (2.67) into (2.66) yields

$$\left\|U_m^{\mathrm{mod}}(\mathcal{T}_{\boldsymbol{u}}\mathrm{X})\right\|_2 = \frac{1}{s'}\left\|F_0\right\|_{L^2}. \tag{2.68}$$

Finally, considering Lemma 2.8 (2.45) concludes the proof. ■

3. From ℂMod to ℝMax.

ℂMod operators are found in ScatterNets and complex-valued convolutional networks [47]. However, they are absent from conventional, freely-trained CNN architectures. Therefore, Theorem 2.9 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 (1.1).

As explained in subsection 1.1, an important number of trained convolution kernels exhibit oscillating patterns with well-defined frequencies and orientations. To elaborate, let $\mathrm{V} \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$ denote such a trained kernel, and consider $\mathrm{W} \in l^2_{\mathbb{C}}(\mathbb{Z}^2)$ as the complex-valued companion of V satisfying (2.6). Then, W has its energy concentrated in a small region of the Fourier domain. We thus emit the hypotheses that $\mathrm{W} \in \mathcal{J}\big(\boldsymbol{\theta},\, \kappa\big)$ (2.5) 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 $U^{\max}_{m,\,q}$ instead of $U^{\max}_{m,\,q}[\mathrm{W}]$, when no ambiguity is possible. In what follows, we establish conditions on W under which ℂMod (1.5) and ℝMax (1.1) 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 $\operatorname{Re} F_1 := F * \operatorname{Re}\overline{\Psi}$, employed in ℝMax, as the real part of the complex-valued wavelet transform output $F_1 := F * \overline{\Psi}$, 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 $\operatorname{Re} F_1$ 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 [9]. This result is hinted in subsection 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 subsection 2.4. However, as explained in subsection 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 subsection 3.4. Then, we provide in subsection 3.5 an upper bound for the expected gap between $\mathbb{C}$Mod and $\mathbb{R}$Max outputs.

3.1. Continuous Framework.

This section, inspired from Waldspurger [50, pp. 190–191], provides an intuition about resemblance between $\mathbb{R}$Max and $\mathbb{C}$Mod in the continuous framework. As will be highlighted in subsection 3.2, adaptation to discrete 2D sequences is not straightforward and will require a probabilistic approach.

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

$$G : (\boldsymbol{x}, \boldsymbol{h}) \mapsto \cos\big(\langle \boldsymbol{\nu}, \boldsymbol{h} \rangle - H(\boldsymbol{x})\big), \tag{3.1}$$

where $H : \mathbb{R}^2 \to [0, 2\pi[$ denotes the phase of $F * \overline{\varPsi}$. Lemma 2.2 introduced low-frequency functions F_0, with slow variations. In a nutshell, since $\operatorname{supp} F_0 \subset B_\infty(\varepsilon/2)$, we can write

$$\|\boldsymbol{h}\|_2 \ll \lambda_{F_0} \implies F_0(\boldsymbol{x} + \boldsymbol{h}) \approx F_0(\boldsymbol{x}), \tag{3.2}$$

where we have defined $\lambda_{F_0} := 2\pi/\varepsilon$. Therefore, according to Proposition 3.1 below, we get the following approximation of $F * \operatorname{Re} \overline{\varPsi}$ in a neighborhood around any point $\boldsymbol{x} \in \mathbb{R}^2$:

$$\|\boldsymbol{h}\|_2 \ll \lambda_{F_0} \implies (F * \operatorname{Re} \overline{\varPsi})(\boldsymbol{x} + \boldsymbol{h}) \approx \big|(F * \overline{\varPsi})(\boldsymbol{x})\big|\, G(\boldsymbol{x}, \boldsymbol{h}). \tag{3.3}$$

Proposition 3.1. *For any* $\boldsymbol{h} \in \mathbb{R}^2$,

$$\Big|(F * \operatorname{Re} \overline{\varPsi})(\boldsymbol{x} + \boldsymbol{h}) - \big|(F * \overline{\varPsi})(\boldsymbol{x})\big|\, G(\boldsymbol{x}, \boldsymbol{h})\Big| \le \big|F_0(\boldsymbol{x} + \boldsymbol{h}) - F_0(\boldsymbol{x})\big|. \tag{3.4}$$

Proof. Let us write:

$$\begin{aligned}
(F * \operatorname{Re} \overline{\varPsi})(\boldsymbol{x} + \boldsymbol{h}) &- \big|(F * \overline{\varPsi})(\boldsymbol{x})\big|\, G(\boldsymbol{x}, \boldsymbol{h}) \\
&= \operatorname{Re}\big((F * \overline{\varPsi})(\boldsymbol{x} + \boldsymbol{h})\big) - \big|(F * \overline{\varPsi})(\boldsymbol{x})\big| \operatorname{Re}\Big(\mathrm{e}^{-i\langle \boldsymbol{\nu}, \boldsymbol{h}\rangle}\, \mathrm{e}^{H(\boldsymbol{x})}\Big) \\
&= \operatorname{Re}\big((F * \overline{\varPsi})(\boldsymbol{x} + \boldsymbol{h})\big) - \operatorname{Re}\Big(\big|(F * \overline{\varPsi})(\boldsymbol{x})\big|\, \mathrm{e}^{H(\boldsymbol{x})}\, \mathrm{e}^{-i\langle \boldsymbol{\nu}, \boldsymbol{h}\rangle}\Big) \\
&= \operatorname{Re}\big((F * \overline{\varPsi})(\boldsymbol{x} + \boldsymbol{h})\big) - \operatorname{Re}\Big((F * \overline{\varPsi})(\boldsymbol{x})\, \mathrm{e}^{-i\langle \boldsymbol{\nu}, \boldsymbol{h}\rangle}\Big) \\
&= \operatorname{Re}\Big((F * \overline{\varPsi})(\boldsymbol{x} + \boldsymbol{h}) - (F * \overline{\varPsi})(\boldsymbol{x})\, \mathrm{e}^{-i\langle \boldsymbol{\nu}, \boldsymbol{h}\rangle}\Big).
\end{aligned}$$

Therefore,

$$\begin{aligned}
\Big|(F * \operatorname{Re} \overline{\varPsi})(\boldsymbol{x} + \boldsymbol{h}) - \big|(F * \overline{\varPsi})(\boldsymbol{x})\big|\, G(\boldsymbol{x}, \boldsymbol{h})\Big| &\le \Big|(F * \overline{\varPsi})(\boldsymbol{x} + \boldsymbol{h}) - (F * \overline{\varPsi})(\boldsymbol{x})\, \mathrm{e}^{-i\langle \boldsymbol{\nu}, \boldsymbol{h}\rangle}\Big| \\
&= \Big|F_0(\boldsymbol{x} + \boldsymbol{h})\, \mathrm{e}^{-i\langle \boldsymbol{\nu}, \boldsymbol{x} + \boldsymbol{h}\rangle} - F_0(\boldsymbol{x})\, \mathrm{e}^{-i\langle \boldsymbol{\nu}, \boldsymbol{x} + \boldsymbol{h}\rangle}\Big|,
\end{aligned}$$

which yields (3.4) and concludes the proof. ■

On the one hand, we consider a continuous equivalent of the $\mathbb{C}$Mod operator $U_m^{\mathrm{mod}}[\mathrm{W}]$ as introduced in (1.5). Such an operator, denoted by $U^{\mathrm{mod}}[\mathit{\Psi}]$, is defined, for any $F \in L^2_{\mathbb{R}}(\mathbb{R}^2)$, by

$$U^{\mathrm{mod}}[\mathit{\Psi}]\,(F) : \boldsymbol{x} \mapsto \left|(F * \overline{\mathit{\Psi}})(\boldsymbol{x})\right|. \tag{3.5}$$

On the other hand, we consider the continuous counterpart of $\mathbb{R}$Max as introduced in (1.1). It is defined as the maximum value of $F * \mathrm{Re}\,\overline{\mathit{\Psi}}$ over a sliding spatial window of size $r > 0$. This is possible because F and $\mathrm{Re}\,\overline{\mathit{\Psi}}$ both belong to $L^2_{\mathbb{R}}(\mathbb{R}^2)$, and therefore $F * \mathrm{Re}\,\overline{\mathit{\Psi}}$ is continuous. Such an operator, denoted by $U_r^{\max}[\mathit{\Psi}]$, is defined, for any $F \in L^2_{\mathbb{R}}(\mathbb{R}^2)$, by

$$U_r^{\max}[\mathit{\Psi}]\,(F) : \boldsymbol{x} \mapsto \max_{\|\boldsymbol{h}\|_\infty \le r} (F * \mathrm{Re}\,\overline{\mathit{\Psi}})(\boldsymbol{x} + \boldsymbol{h}). \tag{3.6}$$

For the sake of conciseness, the parameter between square brackets is ignored from now on. If $r \ll \lambda_{F_0}$, then (3.3) is valid for any $\boldsymbol{h} \in B_\infty(r)$. Then, using (3.5) and (3.6), we get

$$r \ll \lambda_{F_0} \implies U_r^{\max} F(\boldsymbol{x}) \approx U^{\mathrm{mod}} F(\boldsymbol{x}) \max_{\|\boldsymbol{h}\|_\infty \le r} G(\boldsymbol{x}, \boldsymbol{h}). \tag{3.7}$$

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

$$\frac{\pi}{\|\boldsymbol{\nu}\|_2} \le r \ll \frac{2\pi}{\varepsilon} \implies U_r^{\max} F(\boldsymbol{x}) \approx U^{\mathrm{mod}} F(\boldsymbol{x}). \tag{3.8}$$

3.2. Adaptation to Discrete 2D Sequences. As in subsection 2.4, we consider an input image $\mathrm{X} \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$, a complex, analytic convolution kernel $\mathrm{W} \in \mathcal{J}\big(\boldsymbol{\theta}, \kappa\big)$, 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

$$\mathrm{Y}^{\max} := U_{m,\,q}^{\max}[\mathrm{W}]\,(\mathrm{X}) \qquad \text{and} \qquad \mathrm{Y}^{\mathrm{mod}} := U_m^{\mathrm{mod}}[\mathrm{W}]\,(\mathrm{X}), \tag{3.9}$$

where $U_{m,\,q}^{\max}[\mathrm{W}]$ ($\mathbb{R}$Max) and $U_m^{\mathrm{mod}}[\mathrm{W}]$ ($\mathbb{C}$Mod) have been respectively defined in (1.1) and (1.5). As before, in what follows we omit the parameter between square brackets.

We now use the sampling results from Proposition 2.5. Let F_{X} and $\mathit{\Psi}_{\mathrm{W}} \in \mathcal{V}^{(s)}$ denote the functions satisfying (2.23). Recall that the continuous versions of $\mathbb{C}$Mod and $\mathbb{R}$Max operators have been defined in (3.5) and (3.6), respectively. On the one hand, we apply (2.26) with $m \leftarrow 2m$ to $\mathrm{Y}^{\mathrm{mod}}$. For any $\boldsymbol{n} \in \mathbb{Z}^2$,

$$U_m^{\mathrm{mod}} \mathrm{X}[\boldsymbol{n}] = (F_{\mathrm{X}} * \overline{\mathit{\Psi}}_{\mathrm{W}})(\boldsymbol{x_n}) \tag{3.10}$$

$$= U^{\mathrm{mod}} F_{\mathrm{X}}(\boldsymbol{x_n}), \tag{3.11}$$

with $\boldsymbol{x_n} := 2ms\boldsymbol{n}$. On the other hand, we postulate that

$$U_{m,\,q}^{\max} \mathrm{X}[\boldsymbol{n}] = U_r^{\max} F_{\mathrm{X}}(\boldsymbol{x_n}) \tag{3.12}$$

for a certain value of $r \in \mathbb{R}_{>0}$. Then, (3.8) implies $\mathrm{Y}^{\mathrm{mod}} \approx \mathrm{Y}^{\mathrm{max}}$. However, as explained hereafter, (3.12) is not satisfied, due to the discrete nature of the max pooling grid. According to (1.1) and (1.4), we have

$$U^{\max}_{m,\,q}\mathrm{X}[\boldsymbol{n}] = \max_{\|\boldsymbol{p}\|_\infty \le q} \operatorname{Re}\big((\mathrm{X} * \overline{\mathrm{W}}) \downarrow m\big)\,[2\boldsymbol{n} + \boldsymbol{p}]. \tag{3.13}$$

Therefore, according to (2.26) in Proposition 2.5, we get

$$U^{\max}_{m,\,q}\mathrm{X}[\boldsymbol{n}] = \max_{\|\boldsymbol{p}\|_\infty \le q} \left(F_{\mathrm{X}} * \operatorname{Re}\overline{\Psi}_{\mathrm{W}}\right)(\boldsymbol{x_n} + \boldsymbol{h_p})\,, \tag{3.14}$$

with

$$\boldsymbol{x_n} := 2ms\boldsymbol{n} \qquad \text{and} \qquad \boldsymbol{h_p} := ms\boldsymbol{p}. \tag{3.15}$$

By considering $r_q := ms\left(q + \frac{1}{2}\right)$, we get a variant of (3.12) 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 (3.6) with $r \leftarrow r_q$. As a consequence, (3.7) is replaced in the discrete framework by

$$q \ll 2\pi/(m\kappa) \quad \Longrightarrow \quad U^{\max}_{m,\,q}\mathrm{X}[\boldsymbol{n}] \approx U^{\mathrm{mod}}_{m}\mathrm{X}[\boldsymbol{n}] \max_{\|\boldsymbol{p}\|_\infty \le q} G_{\mathrm{X}}\big(\boldsymbol{x_n},\, \boldsymbol{h_p}\big), \tag{3.16}$$

where we have introduced, similarly to (3.1),

$$G_{\mathrm{X}} : (\boldsymbol{x},\, \boldsymbol{h}) \mapsto \cos\big(\langle \boldsymbol{\nu},\, \boldsymbol{h}\rangle - H_{\mathrm{X}}(\boldsymbol{x})\big), \tag{3.17}$$

with

$$\boldsymbol{\nu} := \boldsymbol{\theta}/s \qquad \text{and} \qquad H_{\mathrm{X}} := \angle\left(F_{\mathrm{X}} * \overline{\Psi}_{\mathrm{W}}\right), \tag{3.18}$$

where $\angle : \mathbb{C} \to [0,\, 2\pi[$ denotes the phase operator. Unlike the continuous case, even if the window size r_q is large enough, the existence of $\boldsymbol{p} \in \{-q\,..\,q\}^2$ such that $G_{\mathrm{X}}\big(\boldsymbol{x_n},\, \boldsymbol{h_p}\big) = 1$ is not guaranteed, as illustrated in Figure 3 with $q = 1$. Instead, we can only seek a probabilistic estimation of the normalized mean squared error between $\mathrm{Y}^{\max}$ and $\mathrm{Y}^{\mathrm{mod}}$.

Approximation (3.16) implies

$$q \ll 2\pi/(m\kappa) \implies \left\|U^{\mathrm{mod}}_{m}\mathrm{X} - U^{\max}_{m,\,q}\mathrm{X}\right\|_2 \approx \|\delta_{m,\,q}\mathrm{X}\|_2\,, \tag{3.19}$$

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

$$\delta_{m,\,q}\mathrm{X}[\boldsymbol{n}] := U^{\mathrm{mod}}_{m}\mathrm{X}[\boldsymbol{n}]\left(1 - \max_{\|\boldsymbol{p}\|_\infty \le q} G_{\mathrm{X}}\big(\boldsymbol{x_n},\, \boldsymbol{h_p}\big)\right). \tag{3.20}$$

Expression (3.19) 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 × frequency localization) and the grid half-size q. In what follows, we establish a bound characterizing this approximation, which will be provided in Proposition 3.5.

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

$$A_{\mathrm{X}} : (\boldsymbol{x},\, \boldsymbol{h}) \mapsto (F_{\mathrm{X}} * \operatorname{Re} \overline{\varPsi}_{\mathrm{W}})(\boldsymbol{x} + \boldsymbol{h}); \tag{3.21}$$

$$\widetilde{A}_{\mathrm{X}} : (\boldsymbol{x},\, \boldsymbol{h}) \mapsto \left|(F_{\mathrm{X}} * \overline{\varPsi}_{\mathrm{W}})(\boldsymbol{x})\right| G_{\mathrm{X}}(\boldsymbol{x}, \boldsymbol{h}). \tag{3.22}$$

We now consider, for any $\boldsymbol{n} \in \mathbb{Z}^2$, the vectors $\boldsymbol{h}_{\boldsymbol{n}}^{\max}$ and $\boldsymbol{h'}_{\boldsymbol{n}}^{\max} \in ms\,\{-q\mathinner{..}q\}^2$ achieving the maximum value of $A_{\mathrm{X}}(\boldsymbol{x_n},\, \boldsymbol{h_p})$ and $\widetilde{A}_{\mathrm{X}}(\boldsymbol{x_n},\, \boldsymbol{h_p})$ over the max pooling grid, respectively. They satisfy

$$A_{\mathrm{X}}^{\max}(\boldsymbol{x_n}) := A_{\mathrm{X}}\big(\boldsymbol{x_n},\, \boldsymbol{h}_{\boldsymbol{n}}^{\max}\big) = \max_{\|\boldsymbol{p}\|_\infty \leq q} A_{\mathrm{X}}(\boldsymbol{x_n},\, \boldsymbol{h_p}); \tag{3.23}$$

$$\widetilde{A}_{\mathrm{X}}^{\max}(\boldsymbol{x_n}) := \widetilde{A}_{\mathrm{X}}\big(\boldsymbol{x_n},\, \boldsymbol{h'}_{\boldsymbol{n}}^{\max}\big) = \max_{\|\boldsymbol{p}\|_\infty \leq q} \widetilde{A}_{\mathrm{X}}(\boldsymbol{x_n},\, \boldsymbol{h_p}). \tag{3.24}$$

Then, according to (3.10) and (3.14), we get, for any $\boldsymbol{n} \in \mathbb{Z}^2$,

$$A_{\mathrm{X}}^{\max}(\boldsymbol{x_n}) = U_{m,\,q}^{\max}\mathrm{X}[\boldsymbol{n}]; \tag{3.25}$$

$$\widetilde{A}_{\mathrm{X}}^{\max}(\boldsymbol{x_n}) = U_{m}^{\mathrm{mod}}\mathrm{X}[\boldsymbol{n}] \max_{\|\boldsymbol{p}\|_\infty \leq q} G_{\mathrm{X}}\big(\boldsymbol{x_n},\, \boldsymbol{h_p}\big), \tag{3.26}$$

and (3.16) becomes

$$q \ll 2\pi/(m\kappa) \quad \Longrightarrow \quad A_{\mathrm{X}}^{\max}(\boldsymbol{x_n}) \approx \widetilde{A}_{\mathrm{X}}^{\max}(\boldsymbol{x_n}). \tag{3.27}$$

Remark 3.2. Expression (3.3) implies that, if $q \ll 2\pi/(m\kappa)$, then $A_{\mathrm{X}}(\boldsymbol{x_n},\, \boldsymbol{h_p}) \approx \widetilde{A}_{\mathrm{X}}(\boldsymbol{x_n},\, \boldsymbol{h_p})$ for all $\boldsymbol{p} \in \{-q\mathinner{..}q\}^2$. However, this property does not guarantee that A_{X} and $\widetilde{A}_{\mathrm{X}}$ reach their maximum in the same exact location; *i.e.*, that $\boldsymbol{h}_{\boldsymbol{n}}^{\max} = \boldsymbol{h'}_{\boldsymbol{n}}^{\max}$.

The following lemma provides a bound for approximation (3.27).

Lemma 3.3. *For any $\boldsymbol{x} \in \mathbb{R}^2$,*

$$\left|A_{\mathrm{X}}^{\max}(\boldsymbol{x_n}) - \widetilde{A}_{\mathrm{X}}^{\max}(\boldsymbol{x_n})\right| \leq \max_{\boldsymbol{h} \in \left\{\boldsymbol{h}_{\boldsymbol{n}}^{\max},\, \boldsymbol{h'}_{\boldsymbol{n}}^{\max}\right\}} \left|F_0(\boldsymbol{x_n} + \boldsymbol{h}) - F_0(\boldsymbol{x_n})\right|. \tag{3.28}$$

Proof. We apply Proposition 3.1 with $\boldsymbol{h} \leftarrow \boldsymbol{h}_{\boldsymbol{n}}^{\max}$ and $\boldsymbol{h} \leftarrow \boldsymbol{h'}_{\boldsymbol{n}}^{\max}$, respectively:

$$A_{\mathrm{X}}^{\max}(\boldsymbol{x_n}) \leq \widetilde{A}_{\mathrm{X}}\big(\boldsymbol{x_n},\, \boldsymbol{h}_{\boldsymbol{n}}^{\max}\big) + \left|F_0\big(\boldsymbol{x_n} + \boldsymbol{h}_{\boldsymbol{n}}^{\max}\big) - F_0(\boldsymbol{x_n})\right|; \tag{3.29}$$

$$\widetilde{A}_{\mathrm{X}}^{\max}(\boldsymbol{x_n}) \leq A_{\mathrm{X}}\big(\boldsymbol{x_n},\, \boldsymbol{h'}_{\boldsymbol{n}}^{\max}\big) + \left|F_0\big(\boldsymbol{x_n} + \boldsymbol{h'}_{\boldsymbol{n}}^{\max}\big) - F_0(\boldsymbol{x_n})\right|. \tag{3.30}$$

By construction, we have, for any $\boldsymbol{p} \in \{-q\mathinner{..}q\}^2$,

$$\widetilde{A}_{\mathrm{X}}\big(\boldsymbol{x_n},\, \boldsymbol{h_p}\big) \leq \widetilde{A}_{\mathrm{X}}^{\max}(\boldsymbol{x_n}) \qquad \text{and} \qquad A_{\mathrm{X}}\big(\boldsymbol{x_n},\, \boldsymbol{h_p}\big) \leq A_{\mathrm{X}}^{\max}(\boldsymbol{x_n}). \tag{3.31}$$

Moreover, by definition, there exists $\boldsymbol{p}$ and $\boldsymbol{p}' \in \{-q\mathinner{..}q\}^2$ such that $\boldsymbol{h}_{\boldsymbol{n}}^{\max} = \boldsymbol{h_p}$ and $\boldsymbol{h'}_{\boldsymbol{n}}^{\max} = \boldsymbol{h_{p'}}$. Therefore, (3.29) and (3.30) yield, respectively,

$$A_{\mathrm{X}}^{\max}(\boldsymbol{x_n}) \leq \widetilde{A}_{\mathrm{X}}^{\max}(\boldsymbol{x_n}) + \left|F_0\big(\boldsymbol{x_n} + \boldsymbol{h}_{\boldsymbol{n}}^{\max}\big) - F_0(\boldsymbol{x_n})\right|; \tag{3.32}$$

$$\widetilde{A}_{\mathrm{X}}^{\max}(\boldsymbol{x_n}) \leq A_{\mathrm{X}}^{\max}(\boldsymbol{x_n}) + \left|F_0\big(\boldsymbol{x_n} + \boldsymbol{h'}_{\boldsymbol{n}}^{\max}\big) - F_0(\boldsymbol{x_n})\right|, \tag{3.33}$$

which yields (3.28) and concludes the proof. ■

Before stating Proposition 3.5, we consider the following hypothesis:

Hypothesis 3.4. There exists $\boldsymbol{h}_0 \in \mathbb{R}^2$ with $\|\boldsymbol{h}_0\|_2 = \sqrt{2}qms$, such that

$$\sum_{\boldsymbol{n}\in\mathbb{Z}^2} \max_{\boldsymbol{h}\in\left\{\boldsymbol{h}_{\boldsymbol{n}}^{\max},\, \boldsymbol{h'}_{\boldsymbol{n}}^{\max}\right\}} \Big|F_0(\boldsymbol{x_n}+\boldsymbol{h}) - F_0(\boldsymbol{x_n})\Big|^2 \le \sum_{\boldsymbol{n}\in\mathbb{Z}^2} \Big|F_0(\boldsymbol{x_n}+\boldsymbol{h}_0) - F_0(\boldsymbol{x_n})\Big|^2. \tag{3.34}$$

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

$$\max_{\boldsymbol{h}\in\left\{\boldsymbol{h}_{\boldsymbol{n}}^{\max},\, \boldsymbol{h'}_{\boldsymbol{n}}^{\max}\right\}} \Big|F_0(\boldsymbol{x_n}+\boldsymbol{h}) - F_0(\boldsymbol{x_n})\Big|^2 \le \Big|F_0(\boldsymbol{x_n}+\boldsymbol{h}_0) - F_0(\boldsymbol{x_n})\Big|^2. \tag{3.35}$$

While this might occasionally not be true, Hypothesis 3.4 postulates that, when summing over all the datapoints, the inequality holds.

We now formally state the result characterizing approximation (3.19).

Proposition 3.5. *We assume that condition* (2.53) *is satisfied:* $\kappa \le \pi/m$*. Then, under* Hypothesis 3.4,

$$\big\|U_m^{\mathrm{mod}}\mathrm{X} - U_{m,\,q}^{\max}\mathrm{X}\big\|_2 \le \big\|\delta_{m,\,q}\mathrm{X}\big\|_2 + \beta_q(m\kappa)\,\big\|U_m^{\mathrm{mod}}\mathrm{X}\big\|_2, \tag{3.36}$$

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

$$\beta_q : \kappa' \mapsto q\kappa'. \tag{3.37}$$

Proof. Let us write:

$$\begin{aligned}
\big\|U_m^{\mathrm{mod}}\mathrm{X} - U_{m,\,q}^{\max}\mathrm{X}\big\|_2^2 &= \sum_{\boldsymbol{n}\in\mathbb{Z}^2} \Big(U_m^{\mathrm{mod}}\mathrm{X}[\boldsymbol{n}] - U_{m,\,q}^{\max}\mathrm{X}[\boldsymbol{n}]\Big)^2 \\
&= \sum_{\boldsymbol{n}\in\mathbb{Z}^2} \Bigg(U_m^{\mathrm{mod}}\mathrm{X}[\boldsymbol{n}] - U_m^{\mathrm{mod}}\mathrm{X}[\boldsymbol{n}] \max_{\|\boldsymbol{p}\|_\infty \le q} G_{\mathrm{X}}\big(\boldsymbol{x_n},\, \boldsymbol{h_p}\big) \\
&\qquad\qquad + U_m^{\mathrm{mod}}\mathrm{X}[\boldsymbol{n}] \max_{\|\boldsymbol{p}\|_\infty \le q} G_{\mathrm{X}}\big(\boldsymbol{x_n},\, \boldsymbol{h_p}\big) - U_{m,\,q}^{\max}\mathrm{X}[\boldsymbol{n}]\Bigg)^2 \\
&= \sum_{\boldsymbol{n}\in\mathbb{Z}^2} \Big(\delta_{m,\,q}\mathrm{X}[\boldsymbol{n}] + \widetilde{A}_{\mathrm{X}}^{\max}(\boldsymbol{x_n}) - A_{\mathrm{X}}^{\max}(\boldsymbol{x_n})\Big)^2,
\end{aligned}$$

according to (3.20), (3.25) and (3.26). Then, using the triangle inequality, we get

$$\big\|U_m^{\mathrm{mod}}\mathrm{X} - U_{m,\,q}^{\max}\mathrm{X}\big\|_2 \le \big\|\delta_{m,\,q}\mathrm{X}\big\|_2 + \left(\sum_{\boldsymbol{n}\in\mathbb{Z}^2} \Big(\widetilde{A}_{\mathrm{X}}^{\max}(\boldsymbol{x_n}) - A_{\mathrm{X}}^{\max}(\boldsymbol{x_n})\Big)^2\right)^{1/2}. \tag{3.38}$$

Furthermore, Lemma 3.3 yields

$$\sum_{\boldsymbol{n}\in\mathbb{Z}^2}\left(\widetilde{A}_{\mathrm{X}}^{\max}(\boldsymbol{x_n})-A_{\mathrm{X}}^{\max}(\boldsymbol{x_n})\right)^2\leq\sum_{\boldsymbol{n}\in\mathbb{Z}^2}\max_{\boldsymbol{h}\in\left\{\boldsymbol{h}_{\boldsymbol{n}}^{\max},\,\boldsymbol{h'}_{\boldsymbol{n}}^{\max}\right\}}\left|F_0(\boldsymbol{x_n}+\boldsymbol{h})-F_0(\boldsymbol{x_n})\right|^2 \tag{3.39}$$

$$\leq\sum_{\boldsymbol{n}\in\mathbb{Z}^2}\left|F_0(\boldsymbol{x_n}+\boldsymbol{h}_0)-F_0(\boldsymbol{x_n})\right|^2, \tag{3.40}$$

according to Hypothesis 3.4. Now, since (2.53) is satisfied, we can use Lemma 2.8 (2.44) with $\boldsymbol{h}\leftarrow\boldsymbol{h}_0$. We get

$$\begin{aligned}\sum_{\boldsymbol{n}\in\mathbb{Z}^2}\left(\widetilde{A}_{\mathrm{X}}^{\max}(\boldsymbol{x_n})-A_{\mathrm{X}}^{\max}(\boldsymbol{x_n})\right)^2&\leq\frac{1}{4m^2s^2}\left\|\mathcal{T}_{\boldsymbol{h}_0}F_0-F_0\right\|_{L^2}^2\\&\leq\alpha(\kappa\boldsymbol{h}_0/s)^2\,\frac{1}{4m^2s^2}\left\|F_0\right\|_{L^2}^2\qquad\text{(acc. to Proposition 2.3)}\\&=\alpha(\kappa\boldsymbol{h}_0/s)^2\left\|U_m^{\mathrm{mod}}\mathrm{X}\right\|_2^2.\qquad\text{(acc. to Lemma 2.8 (2.45))}\end{aligned}$$

Since, according to Hypothesis 3.4, $\|\boldsymbol{h}_0\|_2=\sqrt{2}qms$, it comes that $\|\boldsymbol{h}_0\|_1=2qms$. Therefore,

$$\alpha(\kappa\boldsymbol{h}_0/s)^2=\frac{\kappa^2\left\|\boldsymbol{h}_0\right\|_1^2}{4s^2}=(qm\kappa)^2, \tag{3.41}$$

which yields

$$\sum_{\boldsymbol{n}\in\mathbb{Z}^2}\left(\widetilde{A}_{\mathrm{X}}^{\max}(\boldsymbol{x_n})-A_{\mathrm{X}}^{\max}(\boldsymbol{x_n})\right)^2\leq\beta_q(m\kappa)^2\left\|U_m^{\mathrm{mod}}\mathrm{X}\right\|_2^2. \tag{3.42}$$

Finally, plugging (3.42) into (3.38) concludes the proof. ■

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

$$[z,\,z']_{\mathbb{S}^1}:=\left\{z''\in\mathbb{S}^1\;\middle|\;\angle(z^*z'')\leq\angle(z^*z')\right\}. \tag{3.43}$$

We remind readers that $\boldsymbol{x_n}$ and $\boldsymbol{h_p}\in\mathbb{R}^2$ have been defined in (3.15). By using the relation $\cos\alpha=\mathrm{Re}(\mathrm{e}^{i\alpha})$, (3.17) becomes, for any $\boldsymbol{n}\in\mathbb{Z}^2$ and any $\boldsymbol{p}\in\{-q\mathinner{..}q\}^2$,

$$G_{\mathrm{X}}\big(\boldsymbol{x_n},\,\boldsymbol{h_p}\big)=\mathrm{Re}\big(Z_{\mathrm{X}}^*(\boldsymbol{x_n})\,Z_{\boldsymbol{p}}(m\boldsymbol{\theta})\big), \tag{3.44}$$

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

$$Z_{\mathrm{X}}:\boldsymbol{x}\mapsto\mathrm{e}^{i\,H_{\mathrm{X}}(\boldsymbol{x})}\qquad\text{and}\qquad Z_{\boldsymbol{p}}:\boldsymbol{\omega}\mapsto=\mathrm{e}^{i\langle\boldsymbol{\omega},\,\boldsymbol{p}\rangle}, \tag{3.45}$$

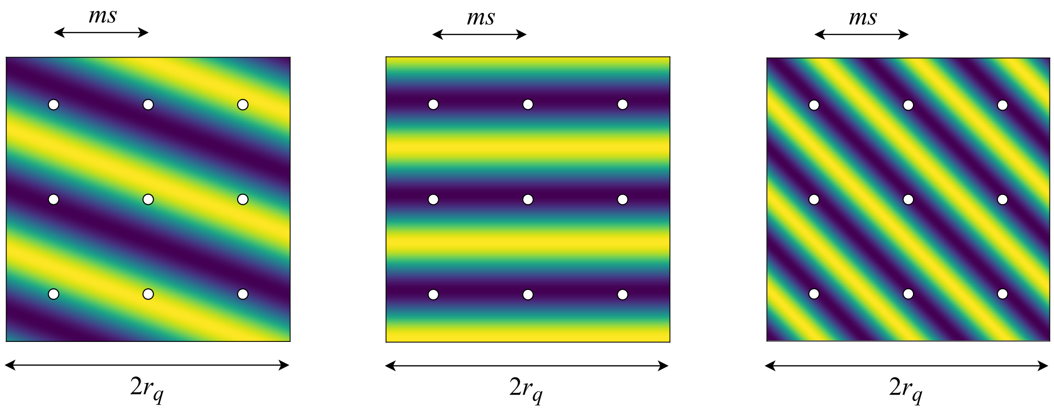

Figure 3. Search for the maximum value of $\boldsymbol{h} \mapsto G_{\mathrm{X}}(\boldsymbol{x},\, \boldsymbol{h})$ over a discrete grid of size 3×3, *i.e.*, $q = 1$. This figure displays 3 examples with different frequencies $\boldsymbol{\nu} := \boldsymbol{\theta}/s$ and phases $H_{\mathrm{X}}(\boldsymbol{x})$. 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 H_{X} denotes the phase of $F_{\mathrm{X}} * \overline{\varPsi}_{\mathrm{W}}$ as introduced in (3.18). On the one hand, $Z_{\mathrm{X}}(\boldsymbol{x_n})$ is the phase (represented on the unit circle $\mathbb{S}^1$) of the complex wavelet transform $F_{\mathrm{X}} * \overline{\varPsi}_{\mathrm{W}}$ at location $\boldsymbol{x_n}$. On the other hand, $Z_{\boldsymbol{p}}(m\boldsymbol{\theta})$ approximates the phase shift between any two evaluations of $F_{\mathrm{X}} * \overline{\varPsi}_{\mathrm{W}}$ at locations $\boldsymbol{x}$, $\boldsymbol{x}'$ such that $\boldsymbol{x}' - \boldsymbol{x} = \boldsymbol{h_p}$. This however is only true if we assume that $\varPsi_{\mathrm{W}}$ exhibits slow amplitude variations. Then, $G_{\mathrm{X}}\big(\boldsymbol{x_n},\, \boldsymbol{h_p}\big)$ approximates the cosine of the phase of $F_{\mathrm{X}} * \overline{\varPsi}_{\mathrm{W}}$ at location $\boldsymbol{x_n} + \boldsymbol{h_p}$.

According to (3.16), $\max_{\|\boldsymbol{p}\|_\infty \leq q} G_{\mathrm{X}}\big(\boldsymbol{x_n},\, \boldsymbol{h_p}\big)$ approximates the ratio between $\mathbb{R}$Max and $\mathbb{C}$Mod outputs at discrete location $\boldsymbol{n} \in \mathbb{Z}^2$. The intuition behind this is that max pooling seeks a point in a discrete grid around $\boldsymbol{x_n}$ where the phase of $F_{\mathrm{X}} * \overline{\varPsi}_{\mathrm{W}}$ is the closest to 1, thereby maximizing the amount of energy on the real part of the signal. Assuming slow amplitude variations of $\varPsi_{\mathrm{W}}$, the result therefore approximates the modulus of the complex coefficients.

To get an estimation of $\delta_{m,\,q}\mathrm{X}[\boldsymbol{n}]$ (3.20), we will exploit the following property. If the phases $Z_{\boldsymbol{p}}(m\boldsymbol{\theta})$ for $\boldsymbol{p} \in \{-q \mathinner{..} q\}^2$ are well distributed on the unit circle, then the values of $G_{\mathrm{X}}\big(\boldsymbol{x_n},\, \boldsymbol{h_p}\big)$ are evenly spread out on $[-1,\, 1]$. Therefore, its maximum value is more likely to be close to 1, and (3.20) becomes

$$\delta_{m,\,q}\mathrm{X}[\boldsymbol{n}] \ll U_m^{\mathrm{mod}}\mathrm{X}[\boldsymbol{n}] \qquad \forall \boldsymbol{n} \in \mathbb{Z}^2. \tag{3.46}$$

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 $\big(Z_i^{(q)}(\boldsymbol{\omega})\big)_{i \in \{0 \mathinner{..} n_q - 1\}}$, obtained by sorting $\{Z_{\boldsymbol{p}}(\boldsymbol{\omega})\}_{\boldsymbol{p} \in \{-q \mathinner{..} q\}^2}$ (3.45) in ascending order of their arguments:

$$0 = H_0^{(q)}(\boldsymbol{\omega}) \leq \cdots \leq H_{n_q-1}^{(q)}(\boldsymbol{\omega}) < 2\pi, \tag{3.47}$$

where $H_i^{(q)}(\boldsymbol{\omega})$ denotes the phase of $Z_i^{(q)}(\boldsymbol{\omega})$. Besides, 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 $Z_i^{(q)}(\boldsymbol{\omega})$:

$$\mathfrak{A}_i^{(q)}(\boldsymbol{\omega}) := \begin{cases} \left[Z_i^{(q)}(\boldsymbol{\omega}),\, Z_{i+1}^{(q)}(\boldsymbol{\omega})\right]_{\mathbb{S}^1} & \text{if } H_{i+1}^{(q)}(\boldsymbol{\omega}) - H_i^{(q)}(\boldsymbol{\omega}) < 2\pi; \\ \mathbb{S}^1 & \text{otherwise.} \end{cases} \tag{3.48}$$

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

$$\delta H_i^{(q)} : \boldsymbol{\omega} \mapsto H_{i+1}^{(q)}(\boldsymbol{\omega}) - H_i^{(q)}(\boldsymbol{\omega}) \tag{3.49}$$

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 discrete 2D stochastic processes. In order to "randomize" F_X introduced in (2.23), we define a continuous stochastic process from X, denoted by F_X, such that

$$\forall \boldsymbol{x} \in \mathbb{R}^2,\ \mathsf{F}_\mathrm{X}(\boldsymbol{x}) := \sum_{\boldsymbol{n} \in \mathbb{Z}^2} \mathrm{X}[\boldsymbol{n}]\, \mathit{\Phi}_{\boldsymbol{n}}^{(s)}(\boldsymbol{x}). \tag{3.50}$$

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

$$\mathsf{M}_\mathrm{X} := |\mathsf{F}_\mathrm{X} * \overline{\mathit{\Psi}}_\mathrm{W}|; \qquad \mathsf{H}_\mathrm{X} := \angle(\mathsf{F}_\mathrm{X} * \overline{\mathit{\Psi}}_\mathrm{W}); \qquad \mathsf{Z}_\mathrm{X} := \mathrm{e}^{i\mathsf{H}_\mathrm{X}}, \tag{3.51}$$

and, for any $\boldsymbol{p} \in \{-q\,..\,q\}^2$,

$$\mathsf{G}_{\mathrm{X},\boldsymbol{p}} := \mathrm{Re}\big(\mathsf{Z}_\mathrm{X}^*\, Z_{\boldsymbol{p}}(m\boldsymbol{\theta})\big); \qquad \mathsf{G}_\mathrm{X}^{\max} := \max_{\|\boldsymbol{p}\|_\infty \le q} \mathsf{G}_{\mathrm{X},\boldsymbol{p}}, \tag{3.52}$$

where the deterministic function $Z_{\boldsymbol{p}}$ has been defined in (3.45).

Remark 3.6. $\mathsf{H}_\mathrm{X}(\boldsymbol{x})$ is ill-defined if $\mathsf{M}_\mathrm{X}(\boldsymbol{x}) = 0$. To overcome this, it is designed to follow a uniform conditional probability distribution on $[0,\, 2\pi[$, given $\mathsf{M}_\mathrm{X}(\boldsymbol{x}) = 0$. Moreover, we impose the following conditional independence, for any $n \in \mathbb{N} \setminus \{0\}$ and $\boldsymbol{x},\, \boldsymbol{y}_0,\, \ldots,\, \boldsymbol{y}_{n-1} \in \mathbb{R}^2$:

$$\mathsf{H}_\mathrm{X}(\boldsymbol{x}) \perp\!\!\!\perp \boldsymbol{\mathsf{M}} \;\big|\; \mathsf{M}_\mathrm{X}(\boldsymbol{x}) = 0, \qquad \text{with} \qquad \boldsymbol{\mathsf{M}} := \big(\mathsf{M}_\mathrm{X}(\boldsymbol{y}_0),\, \cdots,\, \mathsf{M}_\mathrm{X}(\boldsymbol{y}_{n-1})\big)^\top. \tag{3.53}$$

Finally, we impose the following relationship between $\mathsf{H}_{\mathcal{T}_{\boldsymbol{u}}\mathrm{X}}$ and H_X, for any $\boldsymbol{u} \in \mathbb{R}^2$:

$$\mathsf{M}_{\mathcal{T}_{\boldsymbol{u}}\mathrm{X}}(\boldsymbol{x}) = 0 \implies \mathsf{H}_{\mathcal{T}_{\boldsymbol{u}}\mathrm{X}}(\boldsymbol{x}) = \mathcal{T}_{s\boldsymbol{u}}\mathsf{H}_\mathrm{X}(\boldsymbol{x}). \tag{3.54}$$

For any $\boldsymbol{x} \in \mathbb{R}^2$, $F_\mathrm{X}(\boldsymbol{x})$ (2.23) and $H_\mathrm{X}(\boldsymbol{x})$ (3.18) are respectively drawn from $\mathsf{F}_\mathrm{X}(\boldsymbol{x})$ and $\mathsf{H}_\mathrm{X}(\boldsymbol{x})$. Then, $Z_\mathrm{X}(\boldsymbol{x})$ (3.45) is a realization of $\mathsf{Z}_\mathrm{X}(\boldsymbol{x})$. Consequently, according to (3.44), $G_\mathrm{X}\big(\boldsymbol{x},\, \boldsymbol{h}_{\boldsymbol{p}}\big)$ is a realization of $\mathsf{G}_{\mathrm{X},\boldsymbol{p}}(\boldsymbol{x})$. Besides, according to the definition of $\mathbb{C}$Mod in (1.5) and $\boldsymbol{x}_{\boldsymbol{n}}$ in (3.15), Proposition 2.5 with $m \leftarrow 2m$ implies that

$$\mathsf{M}_\mathrm{X}(\boldsymbol{x}_{\boldsymbol{n}}) = U_m^{\mathrm{mod}}\mathrm{X}[\boldsymbol{n}]. \tag{3.55}$$

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 $\mathrm{W} \in \mathcal{J}\big(\boldsymbol{\theta},\, \kappa\big)$. To compute the expected discrepancy between $\mathrm{Y}^{\max}$ and $\mathrm{Y}^{\mathrm{mod}}$, we assume that

$$\|\boldsymbol{\theta}\|_2 \gg 2\pi/N; \tag{3.56}$$

$$\|\boldsymbol{\theta}\|_2 \gg \kappa, \tag{3.57}$$

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 A.

Hypothesis 3.7. For any $\boldsymbol{x} \in \mathbb{R}^2$, $\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x})$ is uniformly distributed on $\mathbb{S}^1$.

Hypothesis 3.8. For any $n \in \mathbb{N} \setminus \{0\}$ and $\boldsymbol{x},\, \boldsymbol{y}_0,\, \ldots,\, \boldsymbol{y}_{n-1} \in \mathbb{R}^2$, the random variables $\mathsf{M}_{\mathrm{X}}(\boldsymbol{y}_i)$ for $i \in \{0\,..\,n-1\}$ are jointly independent of $\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x})$.

3.5. Expected Quadratic Error between $\mathbb{R}$Max and $\mathbb{C}$Mod. In this section, we propose to estimate the expected value of the stochastic quadratic error $\widetilde{\mathsf{P}}_{\mathrm{X}}^2$, defined such that

$$\widetilde{\mathsf{P}}_{\mathrm{X}} := \big\|U_m^{\mathrm{mod}}\mathrm{X} - U_{m,\,q}^{\max}\mathrm{X}\big\|_2 / \big\|U_m^{\mathrm{mod}}\mathrm{X}\big\|_2. \tag{3.58}$$

According to (3.9), this is an estimation of the relative error between $\mathrm{Y}^{\mathrm{mod}}$ and $\mathrm{Y}^{\max}$.

First, let us reformulate $\delta_{m,\,q}\mathrm{X}$, introduced in (3.20), using the probabilistic framework. According to (3.44) and (3.52), we have, for any $\boldsymbol{n} \in \mathbb{Z}^2$,

$$\delta_{m,\,q}\mathrm{X}[\boldsymbol{n}] := U_m^{\mathrm{mod}}\mathrm{X}[\boldsymbol{n}]\,\big(1 - \mathsf{G}_{\mathrm{X}}^{\max}(\boldsymbol{x}_{\boldsymbol{n}})\big). \tag{3.59}$$

We now consider the stochastic process

$$\mathsf{Q}_{\mathrm{X}} := 1 - \mathsf{G}_{\mathrm{X}}^{\max}, \tag{3.60}$$

and the random variable

$$\widetilde{\mathsf{Q}}_{\mathrm{X}} := \|\delta_{m,\,q}\mathrm{X}\|_2 \,/\, \big\|U_m^{\mathrm{mod}}\mathrm{X}\big\|_2. \tag{3.61}$$

The next steps are as follows: (1) at the pixel level, show that $\mathbb{E}[\mathsf{Q}_{\mathrm{X}}(\boldsymbol{x})^2]$ 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 $\widetilde{\mathsf{Q}}_{\mathrm{X}}^2$ is equal to the latter quantity; (3) use Proposition 3.5, which implies that $\widetilde{\mathsf{P}}_{\mathrm{X}} \approx \widetilde{\mathsf{Q}}_{\mathrm{X}}$, to deduce an upper bound on the expected value of $\widetilde{\mathsf{P}}_{\mathrm{X}}^2$.

The first point is established in Proposition 3.9 below, and the two remaining ones are the purpose of Theorem 3.11.

Proposition 3.9. *Assuming Hypothesis 3.7, the expected value of $\mathsf{Q}_{\mathrm{X}}(\boldsymbol{x})^2$ is independent from the choice of $\boldsymbol{x} \in \mathbb{R}^2$, and*

$$\mathbb{E}\left[\mathsf{Q}_{\mathrm{X}}(\boldsymbol{x})^2\right] = \gamma_q(m\boldsymbol{\theta})^2, \tag{3.62}$$

where we have defined

$$\gamma_q : \boldsymbol{\omega} \mapsto \sqrt{\frac{3}{2} + \frac{1}{4\pi} \sum_{i=0}^{n_q-1} \left(\sin \delta H_i^{(q)}(\boldsymbol{\omega}) - 8 \sin \frac{\delta H_i^{(q)}(\boldsymbol{\omega})}{2} \right)}, \tag{3.63}$$

with $\delta H_i^{(q)}(\boldsymbol{\omega}) \in [0,\, 2\pi]$ (3.49) *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 $Z_{\boldsymbol{p}}$ (3.45), $Z_i^{(q)}$, $H_i^{(q)}$ (3.47), $\mathfrak{A}_i^{(q)}$ (3.48), and $\delta H_i^{(q)}$ (3.49); 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 $\left\{[z,\, z']_{\mathbb{S}^1} \mid z,\, z' \in \mathbb{S}^1\right\} \cup \{\mathbb{S}^1\}$, on which we have defined the angular measure ϑ such that $\vartheta(\mathbb{S}^1) := 2\pi$, and

$$\forall z,\, z' \in \mathbb{S}^1,\; \vartheta\left(\left[z,\, z'\right]_{\mathbb{S}^1}\right) := \angle(z^* z'). \tag{3.64}$$

For any $p \in \mathbb{N} \setminus \{0\}$, we compute the p-th moment of $\mathsf{G}_{\mathrm{X}}^{\max}(\boldsymbol{x})$ defined in (3.52). By considering

$$\begin{aligned} g_{\max} : \mathbb{S}^1 &\to [-1,\, 1] \\ z &\mapsto \max_{\|\boldsymbol{p}\|_\infty \le q} \operatorname{Re}\bigl(z^* Z_{\boldsymbol{p}}\bigr), \end{aligned} \tag{3.65}$$

we get $\mathsf{G}_{\mathrm{X}}^{\max}(\boldsymbol{x}) = g_{\max}(\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x}))$. A visual representation of $g_{\max}$ is provided in Figure 4, for two different values of $\boldsymbol{\theta}$. According to Hypothesis 3.7, $\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x})$ follows a uniform distribution on $\mathbb{S}^1$. Therefore,

$$\mathbb{E}\left[\mathsf{G}_{\mathrm{X}}^{\max}(\boldsymbol{x})^p\right] = \frac{1}{2\pi} \int_{\mathbb{S}^1} g_{\max}(z)^p \,\mathrm{d}\vartheta(z), \tag{3.66}$$

which proves that $\mathbb{E}\left[\mathsf{G}_{\mathrm{X}}^{\max}(\boldsymbol{x})^p\right]$ does not depend on $\boldsymbol{x}$. Let us split the unit circle $\mathbb{S}^1$ into the arcs $\mathfrak{A}_0^{(q)}, \ldots, \mathfrak{A}_{n_q-1}^{(q)}$ such as introduced in (3.48):

$$\mathbb{E}\left[\mathsf{G}_{\mathrm{X}}^{\max}(\boldsymbol{x})^p\right] = \frac{1}{2\pi} \sum_{i=0}^{n_q-1} \int_{\mathfrak{A}_i^{(q)}} g_{\max}(z)^p \,\mathrm{d}\vartheta(z). \tag{3.67}$$

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

$$\forall z \in \mathfrak{A}_i^{(q)},\; g_{\max}(z) = \max\Bigl(\operatorname{Re}\bigl(z^* Z_i^{(q)}\bigr),\, \operatorname{Re}\bigl(z^* Z_{i+1}^{(q)}\bigr)\Bigr). \tag{3.68}$$

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

$$\operatorname{Re}\bigl(z^* Z_{i'}^{(q)}\bigr) \le \operatorname{Re}\bigl(z^* Z_i^{(q)}\bigr) \qquad \text{or} \qquad \operatorname{Re}\bigl(z^* Z_{i'}^{(q)}\bigr) \le \operatorname{Re}\bigl(z^* Z_{i+1}^{(q)}\bigr). \tag{3.69}$$

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

$$Z_{i+1}^{(q)} \in \bigl[z,\, Z_{i'}^{(q)}\bigr]_{\mathbb{S}^1}. \tag{3.70}$$

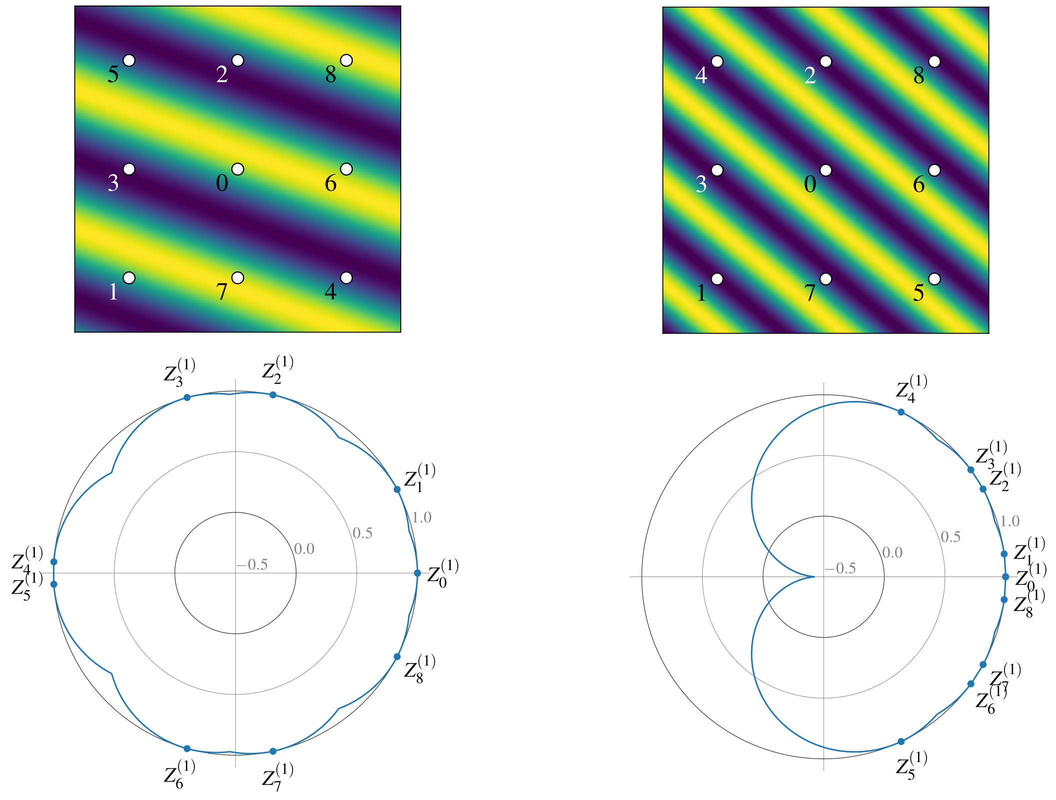

(a) General case

(b) Pathological case

Figure 4. Top: 2D representation of $\boldsymbol{h} \mapsto G_{\mathrm{X}}(\boldsymbol{x_n}, \boldsymbol{h})$ (3.17), for two different values of $\boldsymbol{\theta} \in \mathbb{R}^2$, $q = 1$ and arbitrary values of $m \in \mathbb{N}\backslash\{0\}$ and $s \in \mathbb{R}\backslash\{0\}$. Assuming the plots are centered around $\boldsymbol{h} = \boldsymbol{0}$, each point materializes a location $\boldsymbol{h_p}$ in the max pooling grid, for $\boldsymbol{p} \in \{-q\,..\,q\}^2$. The desirable situation occurs when one of these locations falls near a ridge (bright areas), in which case the outputs produced by $\mathbb{R}$Max and $\mathbb{C}$Mod are similar—see (3.16). Each number $i \in \{0\,..\,8\}$ represents the rank of $Z_{\boldsymbol{p}} \in \mathbb{S}^1$ (3.45), when these values are sorted by ascending order of their arguments (3.47). If rank i is affected to location $\boldsymbol{h_p}$, then we have $Z_{\boldsymbol{p}} = Z_i^{(q)}$. Bottom: polar representations of $g_{\max} : \mathbb{S}^1 \to [-1,\, 1]$ (3.65), corresponding to the same settings. The closer the curve is from the outer ring, the more likely some points $\boldsymbol{h_p}$ will fall near a ridge of G_{X}. (a) Case where the values $Z_{\boldsymbol{p}}$ 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 $Z_{\boldsymbol{p}} = 1$ for any $\boldsymbol{p}$. Figure 3 (middle and right) depicts two such situations.

Therefore, by definition of arcs on the unit circle (3.43), we get

$$\angle\big(z^* Z_{i+1}^{(q)}\big) \leq \angle\big(z^* Z_{i'}^{(q)}\big). \tag{3.71}$$

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

$$\cos\angle\big(z^* Z_{i+1}^{(q)}\big) \geq \cos\angle\big(z^* Z_{i'}^{(q)}\big), \tag{3.72}$$

which yields the right part of (3.69). On the other hand, if $\angle\big(z^* Z_{i'}^{(q)}\big) \geq \pi$, a similar reasoning yields the left part of (3.69). Then, (3.68) holds.

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

$$\overline{Z}_i^{(q)} := \sqrt{Z_i^{(q)} Z_{i+1}^{(q)}}. \tag{3.73}$$

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

$$g_{\max}(z_1) = g_{\max}(z_2). \tag{3.74}$$

A simple calculation yields

$$z_1^* Z_{i+1}^{(q)} = z'^* \widetilde{Z}_i^{(q)} \qquad \text{and} \qquad z_2^* Z_i^{(q)} = \big(z'^* \widetilde{Z}_i^{(q)}\big)^*, \tag{3.75}$$

with

$$\widetilde{Z}_i^{(q)} := \big(Z_i^{(q)*}\, \overline{Z}_i^{(q)}\big) = \big(\overline{Z}_i^{(q)*} Z_{i+1}^{(q)}\big). \tag{3.76}$$

Therefore,

$$\operatorname{Re}\big(z_1^* Z_{i+1}^{(q)}\big) = \operatorname{Re}\big(z_2^* Z_i^{(q)}\big). \tag{3.77}$$

Since z_1, z_2 both belong to $\mathfrak{A}_i^{(q)}$, $g_{\max}(z_1)$ and $g_{\max}(z_2)$ satisfy (3.68). Then, by symmetry, (3.77) implies (3.74). One can observe from Figure 4 that $g_{\max}$ reaches its local minimum at the center of arc $\mathfrak{A}_i^{(q)}$, *i.e.*, $\overline{Z}_i^{(q)}$. This corresponds to a point where $g_{\max}$ is non-differentiable.

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

$$\forall z \in \overline{\mathfrak{A}}_i^{(q)},\ g_{\max}(z) = \operatorname{Re}\big(z^* Z_i^{(q)}\big). \tag{3.78}$$

As a consequence, using symmetry, we get

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

By using the change of variable formula [1, p. 81] with $z \leftarrow \mathrm{e}^{i\eta}$, we get

$$\int_{\mathfrak{A}_i^{(q)}} g_{\max}(z)^p \,\mathrm{d}\vartheta(z) = 2\int_{H_i^{(q)}}^{\overline{H}_i^{(q)}} \cos^p\big(\eta - H_i^{(q)}\big)\,\mathrm{d}\eta, \tag{3.79}$$

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

$$\int_{\mathfrak{A}_i^{(q)}} g_{\max}(z)^p \,\mathrm{d}\vartheta(z) = 2\int_0^{\delta H_i^{(q)}/2} \cos^p \eta' \,\mathrm{d}\eta'. \tag{3.80}$$

Now, we insert (3.80) into (3.67), and compute $\mathbb{E}\left[\mathsf{G}_{\mathrm{X}}^{\max}(\boldsymbol{x})^p\right]$ for $p \leftarrow 1$ and $p \leftarrow 2$:

$$\mathbb{E}\left[\mathsf{G}_{\mathrm{X}}^{\max}(\boldsymbol{x})\right] = \frac{1}{\pi}\sum_{i=0}^{n_q-1} \sin\frac{\delta H_i^{(q)}}{2};$$

$$\mathbb{E}\left[\mathsf{G}_{\mathrm{X}}^{\max}(\boldsymbol{x})^2\right] = \frac{1}{2} + \frac{1}{4\pi}\sum_{i=0}^{n_q-1} \sin\delta H_i^{(q)}.$$

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

$$\mathbb{E}\left[\mathsf{Q}_{\mathrm{X}}(\boldsymbol{x})^2\right] := \frac{3}{2} + \frac{1}{4\pi}\sum_{i=0}^{n_q-1}\left(\sin\delta H_i^{(q)} - 8\sin\frac{\delta H_i^{(q)}}{2}\right), \tag{3.81}$$

which concludes the proof. ■

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

$$\gamma_q(m\boldsymbol{\theta}) = o(1/q^2), \tag{3.82}$$

providing an order-two-polynomial decay rate for $\mathsf{Q}_{\mathrm{X}}(\boldsymbol{x})$, 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 $\big(Z_i^{(q)}(m\boldsymbol{\theta})\big)_{i\in\{0..n_q-1\}}$ is unbalanced.

So far, we established a result at the pixel level. Before stating Theorem 3.11, which extends the result to the image level, we need the following intermediate statement.

Proposition 3.10. *We consider the random variable*

$$\widetilde{\mathsf{S}}_{\mathrm{X}} := \big\|U_m^{\mathrm{mod}}\mathrm{X}\big\|_2. \tag{3.83}$$

Under Hypothesis 3.8, for any $\boldsymbol{x} \in \mathbb{R}^2$,
- *$\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x})$ is independent of $\widetilde{\mathsf{S}}_{\mathrm{X}}$;*
- *$\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x})$, $\mathsf{M}_{\mathrm{X}}(\boldsymbol{x})$ are conditionally independent given $\widetilde{\mathsf{S}}_{\mathrm{X}}$.*

Proof. We suppose that Hypothesis 3.8 is satisfied and we consider $\boldsymbol{x} \in \mathbb{R}^2$. For a given $n \in \mathbb{N}\setminus\{0\}$, we introduce the random variable

$$\widetilde{\mathsf{S}}_{\mathrm{X},\,n} := \sqrt{\sum_{\|\boldsymbol{p}\|_\infty \le n} \mathsf{M}_{\mathrm{X}}(\boldsymbol{x}_{\boldsymbol{p}})^2}. \tag{3.84}$$

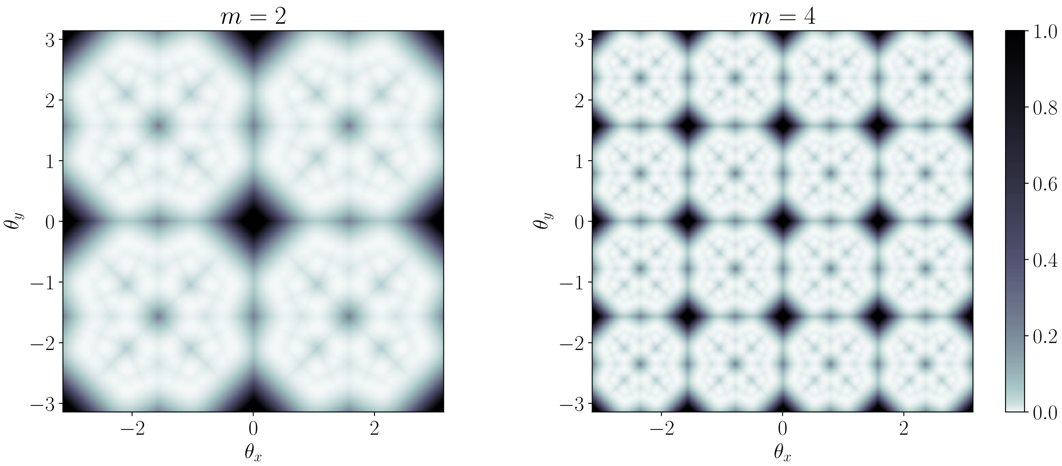

Figure 5. $\gamma(m\boldsymbol{\theta})^2$ as a function of the kernel characteristic frequency $\boldsymbol{\theta} \in [-\pi, \pi]^2$. According to Theorem 3.11, this quantity provides an approximate bound for the expected quadratic error between $\mathbb{R}$Max and $\mathbb{C}$Mod outputs. The subsampling 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 Figure 3 are more likely to occur.

According to Hypothesis 3.8, $\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x})$ is jointly independent of $\mathsf{M}_{\mathrm{X}}(\boldsymbol{x_p})$ for $\boldsymbol{p} \in \{-n \mathinner{..} n\}^2$. Therefore, by composition, $\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x})$ is also independent of $\widetilde{\mathsf{S}}_{\mathrm{X},\,n}$. Moreover, according to (3.55) and (3.83), $\widetilde{\mathsf{S}}_{\mathrm{X},\,n}$ converges almost surely towards $\widetilde{\mathsf{S}}_{\mathrm{X}}$, which proves independence between $\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x})$ and $\widetilde{\mathsf{S}}_{\mathrm{X}}$.

Now, we prove conditional independence between $\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x})$ and $\mathsf{M}_{\mathrm{X}}(\boldsymbol{x})$ given $\widetilde{\mathsf{S}}_{\mathrm{X}}$. According to Hypothesis 3.8,

$$\Big(\mathsf{M}_{\mathrm{X}}(\boldsymbol{x}),\, \widetilde{\mathsf{S}}_{\mathrm{X},\,n}\Big) \perp\!\!\!\perp \mathsf{Z}_{\mathrm{X}}(\boldsymbol{x}), \tag{3.85}$$

where $\perp\!\!\!\perp$ stands for independence. This is because $\widetilde{\mathsf{S}}_{\mathrm{X},\,n}$ only depends on a finite number of $\mathsf{M}_{\mathrm{X}}(\boldsymbol{x_p})$. Therefore,

$$\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x}) \perp\!\!\!\perp \mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) \,\big|\, \widetilde{\mathsf{S}}_{\mathrm{X},\,n}. \tag{3.86}$$

Finally, since $\widetilde{\mathsf{S}}_{\mathrm{X},\,n}$ converges almost surely towards $\widetilde{\mathsf{S}}_{\mathrm{X}}$, it comes that $\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x})$ and $\mathsf{M}_{\mathrm{X}}(\boldsymbol{x})$ are conditionally independent given $\widetilde{\mathsf{S}}_{\mathrm{X}}$. ■

Finally, Propositions 3.9 and 3.10 yield the following theorem. It provides an upper bound on the expected value of the normalized mean squared error $\widetilde{\mathsf{P}}^2_{\mathrm{X}}$, such as defined in (3.58).

Theorem 3.11 (MSE between $\mathbb{C}$Mod and $\mathbb{R}$Max). *Let* $\mathrm{W} \in \mathcal{J}\big(\boldsymbol{\theta},\, \kappa\big)$ *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* $l^2_{\mathbb{R}}(\mathbb{Z}^2)$*. We assume that* *condition* (2.53) *is satisfied:* $\kappa \leq \pi/m$. *Then, under Hypotheses* 3.4, 3.7, *and* 3.8,[2]

$$\mathbb{E}\big[\widetilde{\mathsf{P}}_{\mathrm{X}}^2\big] \leq \big(\beta_q(m\kappa) + \gamma_q(m\boldsymbol{\theta})\big)^2, \tag{3.87}$$

where $\widetilde{\mathsf{P}}_{\mathrm{X}}^2$ (3.58) *denotes the stochastic quadratic error between* $\mathbb{C}$*Mod and* $\mathbb{R}$*Max outputs. We remind that* β_q *and* γ_q *have been introduced in* (3.37) *and* (3.63), *respectively.*

Proof. We consider $\boldsymbol{n} \in \mathbb{Z}^2$. By construction, $\mathsf{Q}_{\mathrm{X}}(\boldsymbol{x_n}) := 1 - \mathsf{G}_{\mathrm{X}}^{\max}(\boldsymbol{x_n})$ only depends on $\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x_n})$. Therefore, under Hypothesis 3.8, Proposition 3.10 implies

$$\mathsf{Q}_{\mathrm{X}}(\boldsymbol{x_n}) \perp\!\!\!\perp \mathsf{M}_{\mathrm{X}}(\boldsymbol{x_n}) \;\big|\; \widetilde{\mathsf{S}}_{\mathrm{X}}^2 \qquad \text{and} \qquad \mathsf{Q}_{\mathrm{X}}(\boldsymbol{x_n}) \perp\!\!\!\perp \widetilde{\mathsf{S}}_{\mathrm{X}}^2. \tag{3.88}$$

Besides, we introduce

$$\widetilde{\Delta}_{\mathrm{X}} := \|\delta_{m,\,q}\mathrm{X}\|_2\,, \tag{3.89}$$

where $\delta_{m,\,q}\mathrm{X}$ is defined in (3.59). Then, using the linearity of $\mathbb{E}$, we get

$$\begin{aligned}
\mathbb{E}\Big[\widetilde{\Delta}_{\mathrm{X}}^2 \;\Big|\; \widetilde{\mathsf{S}}_{\mathrm{X}}^2 = \sigma\Big] &= \sum_{\boldsymbol{n}\in\mathbb{Z}^2} \mathbb{E}\Big[\delta_{m,\,q}[\boldsymbol{n}]^2 \;\Big|\; \widetilde{\mathsf{S}}_{\mathrm{X}}^2 = \sigma\Big] \\
&= \sum_{\boldsymbol{n}\in\mathbb{Z}^2} \mathbb{E}\Big[U_{m,\,l}^{\mathrm{mod}}\mathrm{X}[\boldsymbol{n}]^2\,\big(1 - \mathsf{G}_{\mathrm{X}}^{\max}(\boldsymbol{x_n})\big)^2 \;\Big|\; \widetilde{\mathsf{S}}_{\mathrm{X}}^2 = \sigma\Big] \\
&= \sum_{\boldsymbol{n}\in\mathbb{Z}^2} \mathbb{E}\Big[\mathsf{M}_{\mathrm{X}}(\boldsymbol{x_n})^2\,\mathsf{Q}_{\mathrm{X}}(\boldsymbol{x_n})^2 \;\Big|\; \widetilde{\mathsf{S}}_{\mathrm{X}}^2 = \sigma\Big] \qquad \text{(acc. to (3.55) and (3.60))} \\
&= \sum_{\boldsymbol{n}\in\mathbb{Z}^2} \mathbb{E}\Big[\mathsf{M}_{\mathrm{X}}(\boldsymbol{x_n})^2 \;\Big|\; \widetilde{\mathsf{S}}_{\mathrm{X}}^2 = \sigma\Big]\;\mathbb{E}\big[\mathsf{Q}_{\mathrm{X}}(\boldsymbol{x_n})^2\big] \qquad \text{(acc. to (3.88))}.
\end{aligned}$$

According to (3.55) and (3.83), we have

$$\sum_{\boldsymbol{n}\in\mathbb{Z}^2} \mathsf{M}_{\mathrm{X}}(\boldsymbol{x_n})^2 = \big\|U_m^{\mathrm{mod}}\mathrm{X}\big\|_2^2 = \widetilde{\mathsf{S}}_{\mathrm{X}}^2. \tag{3.90}$$

Therefore, using again the linearity of $\mathbb{E}$, we get

$$\begin{aligned}
\mathbb{E}\Big[\widetilde{\Delta}_{\mathrm{X}}^2 \;\Big|\; \widetilde{\mathsf{S}}_{\mathrm{X}}^2 = \sigma\Big] &= \mathbb{E}\Big[\widetilde{\mathsf{S}}_{\mathrm{X}}^2 \;\Big|\; \widetilde{\mathsf{S}}_{\mathrm{X}}^2 = \sigma\Big]\;\mathbb{E}\big[\mathsf{Q}_{\mathrm{X}}(\boldsymbol{x_n})^2\big] \\
&= \sigma \cdot \mathbb{E}\big[\mathsf{Q}_{\mathrm{X}}(\boldsymbol{x_n})^2\big].
\end{aligned}$$

Under Hypothesis 3.7, Proposition 3.9 yields

$$\mathbb{E}\Big[\widetilde{\Delta}_{\mathrm{X}}^2 \;\Big|\; \widetilde{\mathsf{S}}_{\mathrm{X}}^2 = \sigma\Big] = \sigma \cdot \gamma_q(m\boldsymbol{\theta})^2. \tag{3.91}$$

Besides, we can reformulate $\widetilde{\mathsf{Q}}_{\mathrm{X}}$ such as defined in (3.61): $\widetilde{\mathsf{Q}}_{\mathrm{X}} = \widetilde{\Delta}_{\mathrm{X}}/\widetilde{\mathsf{S}}_{\mathrm{X}}$. Therefore,

$$\mathbb{E}\Big[\widetilde{\mathsf{Q}}_{\mathrm{X}}^2 \;\Big|\; \widetilde{\mathsf{S}}_{\mathrm{X}}^2 = \sigma\Big] = \frac{1}{\sigma}\,\mathbb{E}\Big[\widetilde{\Delta}_{\mathrm{X}}^2 \;\Big|\; \widetilde{\mathsf{S}}_{\mathrm{X}}^2 = \sigma\Big] = \gamma_q(m\boldsymbol{\theta})^2. \tag{3.92}$$

---

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

According to (3.92), the conditional expected value of $\widetilde{\mathsf{Q}}_{\mathrm{X}}^2$ remains the same whatever the outcome of $\widetilde{\mathsf{S}}_{\mathrm{X}}^2$. Thus, the law of total expectation states that

$$\mathbb{E}\big[\widetilde{\mathsf{Q}}_{\mathrm{X}}^2\big] = \mathbb{E}\Big[\mathbb{E}\big[\widetilde{\mathsf{Q}}_{\mathrm{X}}^2 \mid \widetilde{\mathsf{S}}_{\mathrm{X}}^2\big]\Big] = \gamma_q(m\boldsymbol{\theta})^2. \tag{3.93}$$

Since we have assumed Hypothesis 3.4, we can apply Proposition 3.5. Using the definition of $\widetilde{\mathsf{P}}_{\mathrm{X}}$ (3.58) and $\widetilde{\mathsf{Q}}_{\mathrm{X}}$ (3.61), we get

$$\widetilde{\mathsf{P}}_{\mathrm{X}} \leq \widetilde{\mathsf{Q}}_{\mathrm{X}} + \beta_q(m\kappa). \tag{3.94}$$

Then,

$$\mathbb{E}\big[\widetilde{\mathsf{P}}_{\mathrm{X}}^2\big] \leq \mathbb{E}\big[\widetilde{\mathsf{Q}}_{\mathrm{X}}^2\big] + 2\beta_q(m\kappa)\,\mathbb{E}\big[\widetilde{\mathsf{Q}}_{\mathrm{X}}\big] + \beta_q(m\kappa)^2. \tag{3.95}$$

According to Jensen's inequality,

$$\mathbb{E}\big[\widetilde{\mathsf{Q}}_{\mathrm{X}}\big] \leq \sqrt{\mathbb{E}\big[\widetilde{\mathsf{Q}}_{\mathrm{X}}^2\big]} = \gamma_q(m\boldsymbol{\theta}). \tag{3.96}$$

Thus,

$$\mathbb{E}\big[\widetilde{\mathsf{P}}_{\mathrm{X}}^2\big] \leq \gamma_q(m\boldsymbol{\theta})^2 + 2\beta_q(m\kappa)\gamma_q(m\boldsymbol{\theta}) + \beta_q(m\kappa)^2, \tag{3.97}$$

which yields (3.87). ■

Let us analyze the bound obtained in (3.87). 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. Besides, $\beta_q(m\kappa)$ increases linearly with the size of the max pooling grid, which is characterized by q. The second term, $\gamma_q(m\boldsymbol{\theta})$, accounts for the discrete nature of the max pooling grid. It strongly depends on the characteristic frequency $\boldsymbol{\theta}$, as illustrated in Figure 5. According to (3.82), 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 $\mathbb{R}$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 $\mathbb{R}$Max operators. First, we consider the following lemma.

Lemma 4.1. *If Hypotheses 3.7 and 3.8 are satisfied, then they are also true with* $\mathrm{X} \leftarrow \mathcal{T}_{\boldsymbol{u}}\mathrm{X}$, *for any* $\boldsymbol{u} \in \mathbb{R}^2$.

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

$$\mathsf{M}_{\mathcal{T}_{\boldsymbol{u}}\mathrm{X}}(\boldsymbol{x}) = \mathcal{T}_{s\boldsymbol{u}}\mathsf{M}_{\mathrm{X}}(\boldsymbol{x}); \tag{4.1}$$

$$\mathsf{Z}_{\mathcal{T}_{\boldsymbol{u}}\mathrm{X}}(\boldsymbol{x}) = \mathcal{T}_{s\boldsymbol{u}}\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x}). \tag{4.2}$$

According to Lemma 2.6, and since the convolution product commutes with translations, we have

$$\big(F_{\mathcal{T}_{\boldsymbol{u}}\mathrm{X}} * \overline{\varPsi}_{\mathrm{W}}\big)(\boldsymbol{x}) = \mathcal{T}_{s\boldsymbol{u}}\big(F_{\mathrm{X}} * \overline{\varPsi}_{\mathrm{W}}\big)(\boldsymbol{x}). \tag{4.3}$$

Then, using (3.51), the above expression becomes

$$\mathsf{M}_{\mathcal{T}_{\boldsymbol{u}}\mathrm{X}}(\boldsymbol{x}) \times \mathsf{Z}_{\mathcal{T}_{\boldsymbol{u}}\mathrm{X}}(\boldsymbol{x}) = (\mathcal{T}_{s\boldsymbol{u}}\mathsf{M}_{\mathrm{X}})(\boldsymbol{x}) \times (\mathcal{T}_{s\boldsymbol{u}}\mathsf{Z}_{\mathrm{X}})(\boldsymbol{x}). \tag{4.4}$$

Therefore, we necessarily have (4.1). On the one hand, if $\mathsf{M}_{\mathcal{T}_{\boldsymbol{u}}\mathrm{X}}(\boldsymbol{x}) > 0$, then (4.2) is satisfied, by uniqueness of the magnitude-phase decomposition. On the other hand, if $\mathsf{M}_{\mathcal{T}_{\boldsymbol{u}}\mathrm{X}}(\boldsymbol{x}) = 0$, then (3.54) also guarantees (4.2), by design.

Finally, we remind that

$$\mathcal{T}_{s\boldsymbol{u}}\mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) = \mathsf{M}_{\mathrm{X}}(\boldsymbol{x} - s\boldsymbol{u}) \qquad \text{and} \qquad \mathcal{T}_{s\boldsymbol{u}}\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x}) = \mathsf{Z}_{\mathrm{X}}(\boldsymbol{x} - s\boldsymbol{u}). \tag{4.5}$$

Then, considering hypotheses Hypotheses 3.7 and 3.8 with $\boldsymbol{x} \leftarrow \boldsymbol{x} - s\boldsymbol{u}$ yields the result. ■

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

Theorem 4.2 (Shift invariance of $\mathbb{R}$Max). *We assume that the requirements stated in Theorem 3.11 are satisfied. Besides, given a translation vector $\boldsymbol{u} \in \mathbb{R}^2$, we consider the following random variable:*

$$\widetilde{\mathsf{R}}_{\mathrm{X},\,\boldsymbol{u}} := \big\|U^{\max}_{m,\,q}(\mathcal{T}_{\boldsymbol{u}}\mathrm{X}) - U^{\max}_{m,\,q}\mathrm{X}\big\|_2 / \big\|U^{\mathrm{mod}}_{m}\mathrm{X}\big\|_2. \tag{4.6}$$

Then, under condition (2.53), we have

$$\mathbb{E}\big[\widetilde{\mathsf{R}}_{\mathrm{X},\,\boldsymbol{u}}\big] \leq 2\,\big(\beta_q(m\kappa) + \gamma_q(m\boldsymbol{\theta})\big) + \alpha(\kappa\boldsymbol{u}), \tag{4.7}$$

where α, β_q and γ_q are defined in (2.14), (3.37) and (3.63), respectively.

Proof. Using the triangle inequality, we compute

$$\begin{aligned} &\big\|U^{\max}_{m,\,q}(\mathcal{T}_{\boldsymbol{u}}\mathrm{X}) - U^{\max}_{m,\,q}\mathrm{X}\big\|_2 \\ &\qquad \leq \big\|U^{\mathrm{mod}}_{m}(\mathcal{T}_{\boldsymbol{u}}\mathrm{X})\big\|_2\,\widetilde{\mathsf{P}}_{\mathcal{T}_{\boldsymbol{u}}\mathrm{X}} + \big\|U^{\mathrm{mod}}_{m}\mathrm{X}\big\|_2\,\widetilde{\mathsf{P}}_{\mathrm{X}} + \big\|U^{\mathrm{mod}}_{m}(\mathcal{T}_{\boldsymbol{u}}\mathrm{X}) - U^{\mathrm{mod}}_{m}\mathrm{X}\big\|_2, \end{aligned} \tag{4.8}$$

where $\widetilde{\mathsf{P}}_{\mathrm{X}}$ and $\widetilde{\mathsf{P}}_{\mathcal{T}_{\boldsymbol{u}}\mathrm{X}}$ are defined in (3.58). According to (2.53), we can apply Proposition 2.10 on the first term of (4.8):

$$\big\|U^{\mathrm{mod}}_{m}(\mathcal{T}_{\boldsymbol{u}}\mathrm{X})\big\|_2 = \big\|U^{\mathrm{mod}}_{m}\mathrm{X}\big\|_2. \tag{4.9}$$

Moreover, we can apply Theorem 2.9 to the third term of (4.8):

$$\big\|U^{\mathrm{mod}}_{m}(\mathcal{T}_{\boldsymbol{u}}\mathrm{X}) - U^{\mathrm{mod}}_{m}\mathrm{X}\big\|_2 \leq \alpha(\kappa\boldsymbol{u})\,\big\|U^{\mathrm{mod}}_{m}\mathrm{X}\big\|_2. \tag{4.10}$$

We therefore get

$$\big\|U^{\max}_{m,\,q}(\mathcal{T}_{\boldsymbol{u}}\mathrm{X}) - U^{\max}_{m,\,q}\mathrm{X}\big\|_2 \leq \Big[\widetilde{\mathsf{P}}_{\mathcal{T}_{\boldsymbol{u}}\mathrm{X}} + \widetilde{\mathsf{P}}_{\mathrm{X}} + \alpha(\kappa\boldsymbol{u})\Big]\,\big\|U^{\mathrm{mod}}_{m}\mathrm{X}\big\|_2. \tag{4.11}$$

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

$$\mathbb{E}\big[\widetilde{\mathsf{R}}_{\mathrm{X},\,\boldsymbol{u}}\big] \le \mathbb{E}\big[\widetilde{\mathsf{P}}_{\mathcal{T}_{\boldsymbol{u}}\mathrm{X}}\big] + \mathbb{E}\big[\widetilde{\mathsf{P}}_{\mathrm{X}}\big] + \alpha(\kappa\boldsymbol{u}), \tag{4.12}$$

where $\widetilde{\mathsf{R}}_{\mathrm{X},\,\boldsymbol{u}}$ has been introduced in (4.7).

For any stochastic process X' satisfying Hypotheses 3.7 and 3.8, Theorem 3.11 and Jensen's inequality yield:

$$\mathbb{E}\big[\widetilde{\mathsf{P}}_{\mathrm{X}'}\big] \le \beta_q(m\kappa) + \gamma_q(m\boldsymbol{\theta}). \tag{4.13}$$

According to Lemma 4.1, Hypotheses 3.7 and 3.8 are also satisfied for $\mathrm{X} \leftarrow \mathcal{T}_{\boldsymbol{u}}\mathrm{X}$. Therefore, (4.13) is valid for both $\mathrm{X}' \leftarrow \mathrm{X}$ and $\mathrm{X}' \leftarrow \mathcal{T}_{\boldsymbol{u}}\mathrm{X}$, and plugging it into (4.12) concludes the proof. ■

In the bound established in (4.7), the sum $\beta_q(m\kappa) + \gamma_q(m\boldsymbol{\theta})$ accounts for the discrepancy between $\mathbb{R}$Max and $\mathbb{C}$Mod outputs, as stated in Theorem 3.11, whereas the term $\alpha(\kappa\boldsymbol{u})$ characterizes the stability of $\mathbb{C}$Mod outputs, as stated in Theorem 2.9. If κ is sufficiently small, then $\alpha(\kappa\boldsymbol{u})$ and $\beta_q(m\kappa)$ become negligible with respect to $\gamma_q(m\boldsymbol{\theta})$, and the bound can be approximated by $2\,\gamma_q(m\boldsymbol{\theta})$. Theorem 4.2 therefore provides a validity domain for shift invariance of $\mathbb{R}$Max operators, as illustrated in Figure 5 with $q = 1$.

Remark 4.3. The stochastic discrepancy introduced in (4.6) is estimated relatively to the $\mathbb{C}$Mod output. This choice is motivated by the perfect shift invariance of its norm, as shown in Proposition 2.10.

Remark 4.4. 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 2.9, 3.11, and 4.2 to multichannel inputs (*e.g.*, RGB images), employed in conventional CNNs such as AlexNet or ResNet.

First, we define multichannel $\mathbb{R}$Max and $\mathbb{C}$Mod operators relatively to (1.1) and (1.5). We denote by K and $L \in \mathbb{N} \setminus \{0\}$ the number of input and output channels, respectively. Besides, we consider a *multichannel convolution tensor*

$$\mathbf{W} := (\mathrm{W}_{lk})_{l\in\{0..L-1\},\, k\in\{0..K-1\}} \in \big(l^2_{\mathbb{C}}(\mathbb{Z}^2)\big)^{L\times K}. \tag{5.1}$$

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

$$\mathbf{X} := (\mathrm{X}_k)_{k\in\{0..K-1\}} \in \big(l^2_{\mathbb{R}}(\mathbb{Z}^2)\big)^{K}. \tag{5.2}$$

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

$$U^{\max}_{m,\,q,\,l}[\mathbf{W}] : \mathbf{X} \mapsto \operatorname{MaxPool}_q\left(\sum_{k=0}^{K-1} \big(\mathrm{X}_k * \operatorname{Re}\overline{\mathrm{W}}_{lk}\big) \downarrow m\right); \tag{5.3}$$

$$U^{\mathrm{mod}}_{m,\,l}[\mathbf{W}] : \mathbf{X} \mapsto \left|\sum_{k=0}^{K-1} (\mathrm{X}_k * \overline{\mathrm{W}}_{lk}) \downarrow (2m)\right|, \tag{5.4}$$

where $m,\, q \in \mathbb{N} \setminus \{0\}$ respectively denote a subsampling factor and the max pooling grid half-size. Analogously to (3.9) for single-channel inputs, we now consider

$$\mathrm{Y}_l^{\max} := U_{m,\,q,\,l}^{\max}[\mathbf{W}]\,(\mathbf{X}) \qquad \text{and} \qquad \mathrm{Y}_l^{\mathrm{mod}} := U_{m,\,l}^{\mathrm{mod}}[\mathbf{W}]\,(\mathbf{X}). \tag{5.5}$$

Again, in what follows we omit the parameter between square brackets. To apply Theorems 2.9, 3.11, and 4.2 to the current setting on the l-th output channel, we need the following hypotheses.

Hypothesis 5.1 (Monochorome filters). Let

$$\widetilde{\mathrm{W}}_l := \frac{1}{K}\sum_{k=0}^{K-1} \mathrm{W}_{lk} \tag{5.6}$$

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

$$\forall k \in \{0\,..\,K-1\}\,,\; \mathrm{W}_{lk} = \mu_{lk}\widetilde{\mathrm{W}}_l. \tag{5.7}$$

Hypothesis 5.2 (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

$$\widetilde{\mathrm{W}}_l \in \mathcal{J}\big(\boldsymbol{\theta}_l,\, \kappa\big). \tag{5.8}$$

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 5.1, $\mathrm{Y}_l^{\max}$ and $\mathrm{Y}_l^{\mathrm{mod}}$ are the outputs of single-channel $\mathbb{R}$Max and $\mathbb{C}$Mod operators, as introduced in (1.1) and (1.5):

$$\mathrm{Y}_l^{\max} = U_{m,\,q}^{\max}\big[\widetilde{\mathrm{W}}_l\big]\big(\mathrm{X}_l^{\mathrm{lum}}\big) \qquad \text{and} \qquad \mathrm{Y}_l^{\mathrm{mod}} = U_m^{\mathrm{mod}}\big[\widetilde{\mathrm{W}}_l\big]\big(\mathrm{X}_l^{\mathrm{lum}}\big), \tag{5.9}$$

where $\mathrm{X}_l^{\mathrm{lum}} \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$ ("luminance" image) is defined as the following linear combination:

$$\mathrm{X}_l^{\mathrm{lum}} := \sum_{k=0}^{K-1} \mu_{lk}\mathrm{X}_k. \tag{5.10}$$

The results established for single-channel inputs can therefore be extended to multichannel operators. Specifically, we get the following corollaries to Theorems 2.9, 3.11, and 4.2.

Corollary 5.3 (Shift invariance of $\mathbb{C}$Mod). *For a given output channel $l \in \{0\,..\,L-1\}$, we postulate Hypotheses 5.1 and 5.2. Then, for any input image $\mathbf{X} \in \big(l^2_{\mathbb{R}}(\mathbb{Z}^2)\big)^K$ with finite support and any translation vector $\boldsymbol{u} \in \mathbb{R}^2$,*

$$\big\|U_{m,\,l}^{\mathrm{mod}}(\mathcal{T}_{\boldsymbol{u}}\mathbf{X}) - U_{m,\,l}^{\mathrm{mod}}\mathbf{X}\big\|_2 \leq \alpha(\kappa\boldsymbol{u})\,\big\|U_{m,\,l}^{\mathrm{mod}}\mathbf{X}\big\|_2, \tag{5.11}$$

where α has been defined in (2.14).

Corollary 5.4 (MSE between $\mathbb{C}$Mod and $\mathbb{R}$Max). *As in Corollary 5.3, we postulate Hypotheses 5.1 and 5.2. Again, we assume that condition (2.53) is satisfied: $\kappa \leq \pi/m$. Besides, we consider $\mathbf{X}$ as a stack of K discrete stochastic processes, and assume Hypotheses 3.4, 3.7, and 3.8 with $\mathrm{X} \leftarrow \mathrm{X}_l^{\mathrm{lum}}$ and $\mathrm{W} \leftarrow \widetilde{\mathrm{W}}_l$. Then,*

$$\mathbb{E}\big[\widetilde{\mathsf{P}}^2_{\mathbf{X},\,l}\big] \leq \big(\beta_q(m\kappa) + \gamma_q(m\boldsymbol{\theta}_l)\big)^2, \tag{5.12}$$

where we have defined the following random variable:

$$\widetilde{\mathsf{P}}_{\mathbf{X},\,l} := \big\|U^{\mathrm{mod}}_{m,\,l}\mathbf{X} - U^{\mathrm{max}}_{m,\,l}\mathbf{X}\big\|_2 / \big\|U^{\mathrm{mod}}_{m,\,l}\mathbf{X}\big\|_2. \tag{5.13}$$

Corollary 5.5 (Shift invariance of $\mathbb{R}$Max). *We assume that the requirements stated in Corollary 5.4 are satisfied. Then, for any translation vector $\boldsymbol{u} \in \mathbb{R}^2$,*

$$\mathbb{E}\big[\widetilde{\mathsf{R}}_{\mathbf{X},\,\boldsymbol{u},\,l}\big] \leq 2\,\big(\beta_q(m\kappa) + \gamma_q(m\boldsymbol{\theta}_l)\big) + \alpha(\kappa\boldsymbol{u}), \tag{5.14}$$

where we have defined the following random variable:

$$\widetilde{\mathsf{R}}_{\mathbf{X},\,\boldsymbol{u},\,l} := \big\|U^{\mathrm{max}}_{m,\,l}(\mathcal{T}_{\boldsymbol{u}}\mathbf{X}) - U^{\mathrm{max}}_{m,\,l}\mathbf{X}\big\|_2 / \big\|U^{\mathrm{mod}}_{m,\,l}\mathbf{X}\big\|_2. \tag{5.15}$$

Remark 5.6. In the above results, we used a translation operator on multichannel tensors, obtained by applying $\mathcal{T}_{\boldsymbol{u}}$, as defined in (2.34), to each channel 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 2.9, 3.11, and 4.2. To this end, we consider a fully-deterministic scenario implementing the dual-tree complex wavelet packet transform (DT-$\mathbb{C}$WPT), 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 subsection 6.1, DT-$\mathbb{C}$WPT 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 subsection 6.2, we build $\mathbb{C}$Mod and $\mathbb{R}$Max operators based on DT-$\mathbb{C}$WPT convolution kernels.

6.1. Main Properties.

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

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

$$\kappa_J := \pi/m_J, \qquad \text{with} \qquad m_J := 2^{J-1}. \tag{6.1}$$

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

$$\mathbf{D}^{(J)} := \big(\mathrm{D}_l^{\nearrow(J)},\, \mathrm{D}_l^{\searrow(J)},\, \mathrm{D}_l^{\swarrow(J)},\, \mathrm{D}_l^{\nwarrow(J)}\big)_{l\in\{0..4^J-1\}}, \tag{6.2}$$

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

$$\mathrm{D}_l^{\nearrow(J)} = \left(\mathrm{X} * \overline{\mathrm{W}}_l^{\nearrow(J)*}\right) \downarrow 2^J. \tag{6.3}$$

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

$$\mathrm{W}_l^{\nearrow(J)} \in \mathcal{J}\big(\boldsymbol{\theta}_l^{(J)},\, \kappa_J\big) \tag{6.4}$$

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 (2.5). 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 $\mathbf{W}^{\nearrow(J)} := \big(\mathrm{W}_l^{\nearrow(J)}\big)_{l\in\{0..4^J-1\}}$ and $\mathbf{W}^{\searrow(J)} := \big(\mathrm{W}_l^{\searrow(J)}\big)_{l\in\{0..4^J-1\}}$ are provided in Figure 6 with $J = 2$ (Figure 6a, 32 filters) and $J = 3$ (Figure 6b, 128 filters).

The $\mathbb{R}$Max and $\mathbb{C}$Mod operators implemented in our experiments respectively satisfy (1.1) and (1.5) with with $\mathrm{W} \leftarrow \mathrm{W}_l^{\nearrow(J)}$ or $\mathrm{W}_l^{\searrow(J)}$, 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 (2.53) $\kappa \leftarrow \kappa_J$ and $m \leftarrow m_J$.

Remark 6.1. Because X is real-valued, the feature maps $\mathrm{D}_l^{\swarrow(J)}$ and $\mathrm{D}_l^{\nwarrow(J)}$ are the respective complex conjugates of $\mathrm{D}_l^{\nearrow(J)}$ and $\mathrm{D}_l^{\searrow(J)}$, and thus do not need to be explicitly computed. Then, we can easily show that $\mathrm{W}_l^{\swarrow(J)}$ and $\mathrm{W}_l^{\nwarrow(J)}$ are also the complex conjugates of $\mathrm{W}_l^{\nearrow(J)}$ and $\mathrm{W}_l^{\searrow(J)}$, respectively.

6.2. DT-$\mathbb{C}$WPT-Based $\mathbb{R}$Max and $\mathbb{C}$Mod Operators. According to (6.1), (6.3), and (6.4), we can apply Theorems 2.9, 3.11, and 4.2 to the dual-tree framework. More precisely, for any output channel $l \in \{0\,.\,.\,4^J - 1\}$, we consider the following $\mathbb{R}$Max and $\mathbb{C}$Mod operators:

$$U_l^{\max\nearrow} : \mathrm{X} \mapsto \operatorname{MaxPool}\left(\left(\mathrm{X} * \operatorname{Re}\overline{\mathrm{W}}_l^{\nearrow(J)}\right) \downarrow 2^{J-1}\right); \tag{6.5}$$

$$U_l^{\mathrm{mod}\nearrow} : \mathrm{X} \mapsto \left|\left(\mathrm{X} * \overline{\mathrm{W}}_l^{\nearrow(J)}\right) \downarrow 2^J\right|. \tag{6.6}$$

Using the notations introduced in (1.5) and (1.1), we have

$$U_l^{\max\nearrow} = U_{m_J}^{\max}\big[\mathrm{W}_l^{\nearrow(J)}\big] \qquad \text{and} \qquad U_l^{\mathrm{mod}\nearrow} := U_{m_J}^{\mathrm{mod}}\big[\mathrm{W}_l^{\nearrow(J)}\big], \tag{6.7}$$

where we have defined $m_J := 2^{J-1}$. Note that, following Remark 4.4, 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 6.2. Both $U_l^{\max\nearrow}$ and $U_l^{\mathrm{mod}\nearrow}$ are implemented using DT-$\mathbb{C}$WPT with J decomposition stages. However, in (6.5), the subsampling factor is equal to 2^{J-1}, instead of 2^J, as stated in (6.3). In order to accommodate this property of $\mathbb{R}$Max operators, the last stage of

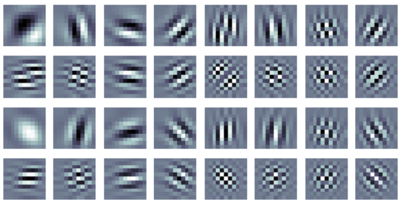

(a) $J = 2$

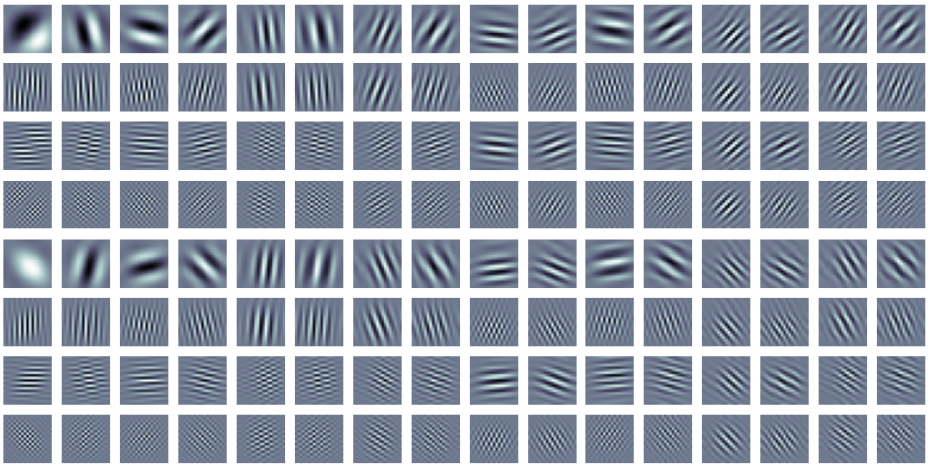

(b) $J = 3$

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 [19]. 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.

DT-$\mathbb{C}$WPT decomposition is carried out without subsampling, resulting in higher redundancy. This is similar to the concept of stationary wavelet transform as described by Nason and Silverman [31]. Furthermore, only the real component of the wavelet feature maps is preserved. On the other hand, $U_l^{\mathrm{mod}\nearrow}$ 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 $\mathbb{R}$Max and $\mathbb{C}$Mod operators $U_l^{\max\nearrow}$ and $U_l^{\mathrm{mod}\nearrow}$, as introduced in (6.5) and (6.6), with both $J = 2$ and 3 stages of wavelet packet decomposition. To cover the whole frequency plane, we also implemented similar operators, denoted by $U_l^{\max\searrow}$ and $U_l^{\mathrm{mod}\searrow}$. They are associated with the convolution filters $\mathrm{W}_l^{\searrow(J)}$, introduced in Proposition B.2, with energy being located in the bottom-right quadrant. However, as explained in Remark 6.1, we did not need to deal with the two other quadrants (negative x-values). Using the validation set of ImageNet-1K [39], ($N := 50\,000$ images), we measured the mean discrepancy between $\mathbb{R}$Max and $\mathbb{C}$Mod outputs, and evaluated the shift invariance

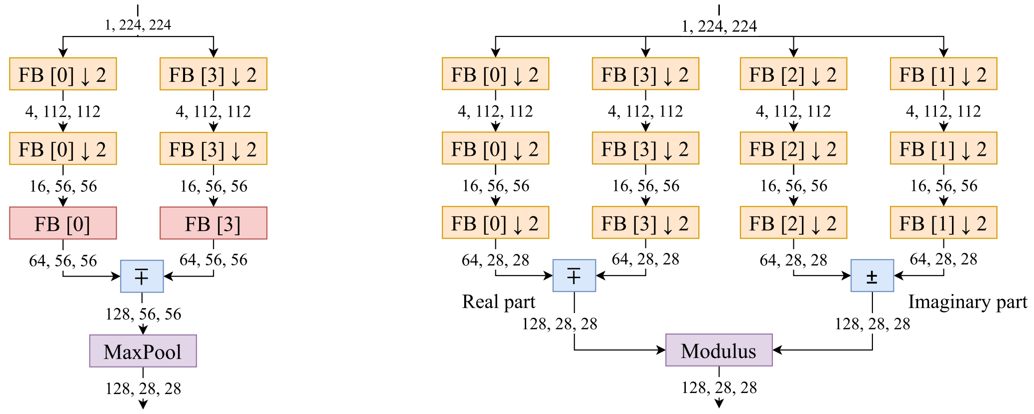

(a) $\mathbb{R}$Max (b) $\mathbb{C}$Mod

Figure 7. Detailed illustration of the $\mathbb{R}$Max (a) and $\mathbb{C}$Mod (b) operators based on DT-$\mathbb{C}$WPT, 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 (B.1). The FB index is indicated between square brackets. The $\mathbb{R}$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 (B.6).

of both models. Dual-tree decompositions have been performed with Q-shift orthogonal filters of length 10 [19].

6.3.1. MSE between $\mathbb{R}$Max and $\mathbb{C}$Mod. Each image $n \in \{0 \,..\, N-1\}$ in the dataset was converted to grayscale, from which a center crop of size 224×224 was extracted. We denote by $\mathrm{X}_n \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$ the resulting input feature map. For any $l \in \big\{0 \,..\, 4^J - 1\big\}$, we denote by

$$\mathrm{Y}_{nl}^{\max\nearrow} := U_l^{\max\nearrow}(\mathrm{X}_n) \qquad \text{and} \qquad \mathrm{Y}_{nl}^{\mathrm{mod}\nearrow} := U_l^{\mathrm{mod}\nearrow}(\mathrm{X}_n) \tag{6.8}$$

the outputs of the l-th $\mathbb{R}$Max and $\mathbb{C}$Mod operators as defined in (6.5) and (6.6), respectively. We adopt similar notations for the bottom-right Fourier quadrant. Then, the normalized mean squared error between $\mathrm{Y}_{nl}^{\mathrm{mod}\nearrow}$ and $\mathrm{Y}_{nl}^{\max\nearrow}$ was computed. It is defined by the square of

$$\rho_{nl}^{\nearrow} := \big\|\mathrm{Y}_{nl}^{\mathrm{mod}\nearrow} - \mathrm{Y}_{nl}^{\max\nearrow}\big\|_2 \big/ \big\|\mathrm{Y}_{nl}^{\mathrm{mod}\nearrow}\big\|_2. \tag{6.9}$$

Finally, the for each output channel l, an empirical estimate for $\mathbb{E}\big[\widetilde{\mathsf{P}}_{\mathrm{X}}^2\big]$, introduced in (3.58), was obtained by averaging $\rho_{nl}^{\nearrow 2}$ over the whole dataset. We denote by $\widetilde{\rho}_l^{\nearrow 2}$ the corresponding quantity.

Since $U_l^{\max\nearrow}$ and $U_l^{\mathrm{mod}\nearrow}$ are parameterized by $\mathrm{W}_l^{\nearrow(J)}$, it follows that $\widetilde{\rho}_l^{\nearrow 2}$ depends on the filter's characteristic frequency $\boldsymbol{\theta}_l^{(J)}$ (6.4). According to Proposition B.4, these frequencies form a regular grid in the top-right quadrant of Fourier domain. This provides a visual representation of $\widetilde{\rho}_l^{\nearrow 2}$, as shown in Figure 8. This figure also displays $\widetilde{\rho}_l^{\searrow 2}$, corresponding to

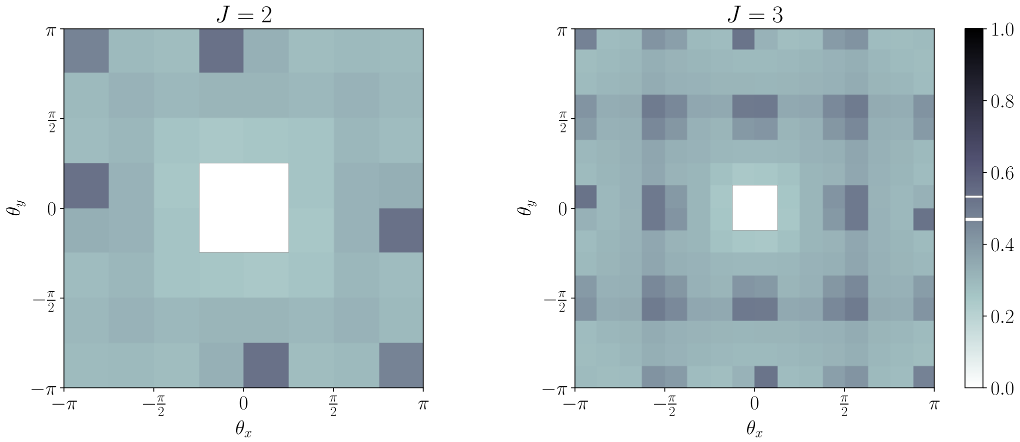

Figure 8. Empirical estimates of the normalized mean squared error between $\mathbb{R}$Max and $\mathbb{C}$Mod outputs, computed on ImageNet-1K (validation set). For each channel $l \in \left\{0\,..\,4^J - 1\right\}$, $\widetilde{\rho}_l^{\nearrow 2}$ is plotted as a grayscale pixel centered in $\boldsymbol{\theta}_l^{(J)}$ such as introduced in (6.4) (top-right quadrant). Similarly, $\widetilde{\rho}_l^{\searrow 2}$ is plotted in the bottom-right quadrant. Finally, the bottom- and top-left quadrants ($\widetilde{\rho}_l^{\swarrow 2}$ and $\widetilde{\rho}_l^{\nwarrow 2}$) 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 Figure 5. Note that the low-pass filters have been discarded because they are outside the scope of this study.

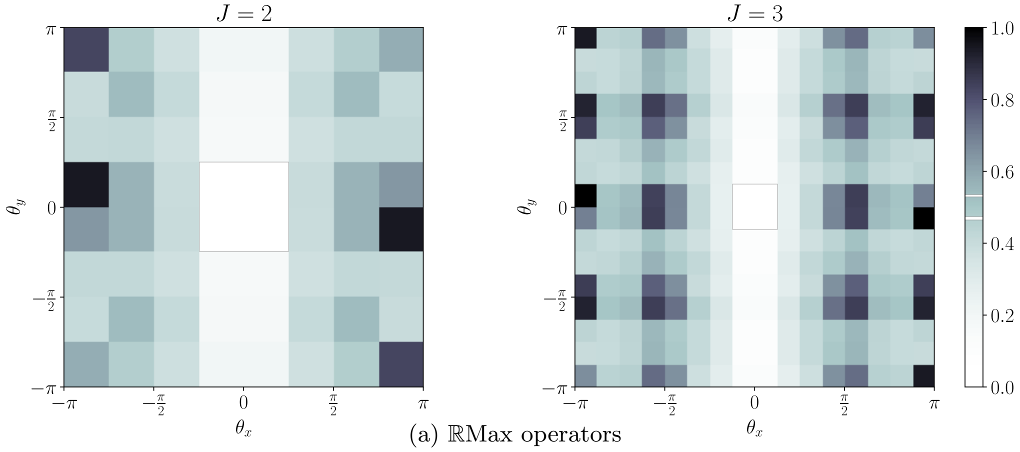

(a) $\mathbb{R}$Max operators

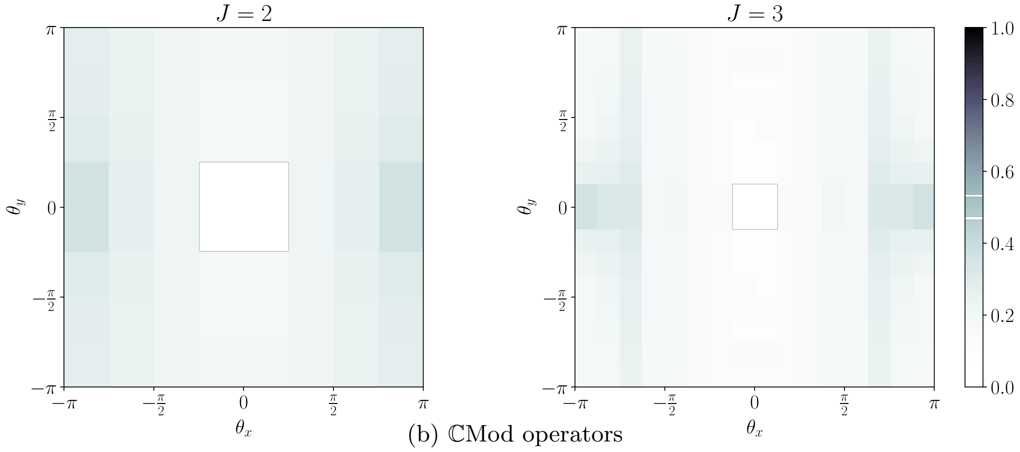

(b) $\mathbb{C}$Mod operators

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

the bottom-right quadrant. The half-plane of negative x-values has simply been symmetrized, following Remark 6.1. 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 $\mathrm{W}_l^{\nearrow(J)}$ or $\mathrm{W}_l^{\searrow(J)}$ overlaps a dark region of Figure 5. This result corroborates Theorem 4.2, 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:

$$\rho_{nl}^{\max\nearrow} := \big\|\mathrm{Y}_{nl}^{\max'\nearrow} - \mathrm{Y}_{nl}^{\max\nearrow}\big\|_2 \,/\, \big\|\mathrm{Y}_{nl}^{\mathrm{mod}\nearrow}\big\|_2, \tag{6.10}$$

where $\mathrm{Y}_{nl}^{\max'\nearrow}$ satisfies (6.7) with $\mathrm{X}_n \leftarrow \mathrm{X}'_n$. Finally, for each output channel $l \in \{0\,..\,4^J - 1\}$, an empirical estimate for $\mathbb{E}\big[\widetilde{\mathsf{R}}_{\mathrm{X},\boldsymbol{u}}\big]$, satisfying (4.6) with $\boldsymbol{u} = (1,\, 0)^\top$, was obtained by averaging $\rho_{nl}^{\max\nearrow}$ over the whole dataset. We denote by $\widetilde{\rho}_l^{\max\nearrow}$ the corresponding quantity. We point out that shift invariance is measured relatively to the norm of the $\mathbb{C}$Mod output, as explained in Remark 4.3.

On the other hand, the same procedure was applied to the $\mathbb{C}$Mod operators:

$$\rho_{nl}^{\mathrm{mod}\nearrow} := \big\|\mathrm{Y}_{nl}^{\mathrm{mod}'\nearrow} - \mathrm{Y}_{nl}^{\mathrm{mod}\nearrow}\big\|_2 \,/\, \big\|\mathrm{Y}_{nl}^{\mathrm{mod}\nearrow}\big\|_2, \tag{6.11}$$

and $\widetilde{\rho}_l^{\mathrm{mod}\nearrow}$ was obtained as before by averaging $\rho_{nl}^{\mathrm{mod}\nearrow}$ over the whole dataset.

A visual representation of $\widetilde{\rho}_l^{\max\nearrow}$ and $\widetilde{\rho}_l^{\mathrm{mod}\nearrow}$ are provided in Figure 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 $\mathbb{R}$Max and $\mathbb{C}$Mod outputs (Figure 8) are correlated with shift instability of $\mathbb{R}$Max (Figure 9, top). This is in line with (3.87) and (4.7) in Theorems 3.11 and 4.2. Note that $\mathbb{C}$Mod outputs are nearly shift invariant regardless the characteristic frequency $\boldsymbol{\theta}_l^{(J)}$ (Figure 9, bottom), as predicted by Theorem 2.9 (2.54).

7. Conclusion. In this paper, we explored the shift invariance properties captured by the max pooling operator, when applied on top of 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 $\mathbb{C}$Mod operator can serve as a stable proxy for $\mathbb{R}$Max, extracting comparable features, except for certain filter frequencies, for which potential degeneracies can arise after max pooling. This suggests a promising approach for improving shift invariance in CNNs while preserving high-frequency information. This is the main focus of [25], in which we apply these principles to real-life architectures.

A link was missing between real- and complex-valued convolutions in CNNs. By comparing the outputs of $\mathbb{C}$Mod and $\mathbb{R}$Max 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. [47], the correspondence between standard real-valued CNNs (using max pooling) and complex wavelets is no longer "just a vague analogy."

Appendix A. Theoretical Foundations for our Hypotheses.

In this section, we provide theoretical arguments for justifying Hypotheses 3.7 and 3.8. Given $n \in \mathbb{N} \setminus \{0\}$, we define *n-th order stationarity* of a given stochastic process F as stated by Park et al. [34, p. 152]: for any $n' \in \{0 \,..\, n-1\}$, $(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_{n'}) \in (\mathbb{R}^2)^{n'}$ and $\boldsymbol{h} \in \mathbb{R}^2$, the joint distribution of $\big(\mathsf{F}(\boldsymbol{x}_1), \ldots, \mathsf{F}(\boldsymbol{x}_{n'})\big)$ is identical to the one of $\big(\mathsf{F}(\boldsymbol{x}_1 + \boldsymbol{h}), \ldots, \mathsf{F}(\boldsymbol{x}_{n'} + \boldsymbol{h})\big)$. Besides, *strict-sense stationarity* is defined as n-th order stationarity for any $n \in \mathbb{N} \setminus \{0\}$.

We recall that $\boldsymbol{\nu} := \boldsymbol{\theta}/s$. We then state the following results.

Proposition A.1. *We assume that F_{X} is first-order stationary. If, for any $\boldsymbol{x} \in \mathbb{R}^2$ and any $\boldsymbol{h} \in B_2(2\pi/\left\|\boldsymbol{\nu}\right\|_2)$,*

$$(\mathcal{T}_{\boldsymbol{h}}\mathsf{F}_{\mathrm{X}} * \overline{\Psi}_{\mathrm{W}})(\boldsymbol{x}) = \mathrm{e}^{i\langle \boldsymbol{\nu},\, \boldsymbol{h}\rangle}(\mathsf{F}_{\mathrm{X}} * \overline{\Psi}_{\mathrm{W}})(\boldsymbol{x}), \tag{A.1}$$

then Hypothesis 3.7 is satisfied.

Proof. Let $\boldsymbol{x} \in \mathbb{R}^2$. By design (see Remark 3.6), $\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x})$ follows a uniform conditional probability distribution on $\mathbb{S}^1$, given $\mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) = 0$. In any other cases, we show that the conditional probability measure of $\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x})$ given $\mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) > 0$ is invariant with respect to phase shifts, and is therefore equal to the uniform probability measure on $\mathbb{S}^1$. Specifically, we show that, for any measurable set $\mathfrak{A} \subset \mathbb{S}^1$,

$$\forall \omega \in [0,\, 2\pi]\,,\, \mu(\mathfrak{A}) = \mu(\mathrm{e}^{i\omega}\mathfrak{A}), \tag{A.2}$$

where we have denoted

$$\mu : \mathfrak{A} \mapsto \mathbb{P}\left\{\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x}) \in \mathfrak{A} \mid \mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) > 0\right\}. \tag{A.3}$$

Let $\boldsymbol{h} \in B_2(2\pi/\left\|\boldsymbol{\nu}\right\|_2)$. According to (A.1), and assuming $\mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) > 0$, we get

$$\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x}) \in \mathfrak{A} \iff \mathcal{T}_{\boldsymbol{h}}\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x}) \in \mathrm{e}^{i\langle \boldsymbol{\nu},\, \boldsymbol{h}\rangle}\mathfrak{A}. \tag{A.4}$$

Therefore,

$$\mathbb{P}\left\{\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x}) \in \mathfrak{A} \mid \mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) > 0\right\} = \mathbb{P}\Big\{\mathcal{T}_{\boldsymbol{h}}\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x}) \in \mathrm{e}^{i\langle \boldsymbol{\nu},\, \boldsymbol{h}\rangle}\mathfrak{A} \;\Big|\; \mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) > 0\Big\}. \tag{A.5}$$

Since F_{X} is first-order stationary, $\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x})$ and $\mathcal{T}_{\boldsymbol{h}}\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x})$ have the same conditional probability distribution given $\mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) > 0$. Thus we get

$$\mathbb{P}\left\{\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x}) \in \mathfrak{A} \mid \mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) > 0\right\} = \mathbb{P}\Big\{\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x}) \in \mathrm{e}^{i\langle \boldsymbol{\nu},\, \boldsymbol{h}\rangle}\mathfrak{A} \;\Big|\; \mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) > 0\Big\}. \tag{A.6}$$

Let $\omega \in [0,\, 2\pi]$. Considering $\boldsymbol{h} := \omega\, \boldsymbol{\nu}/\left\|\boldsymbol{\nu}\right\|_2^2$, we have

$$\boldsymbol{h} \in B_2(2\pi/\left\|\boldsymbol{\nu}\right\|_2) \qquad \text{and} \qquad \langle \boldsymbol{\nu},\, \boldsymbol{h}\rangle = \omega. \tag{A.7}$$

Therefore,

$$\forall \omega \in [0,\, 2\pi]\,,\, \mathbb{P}\left\{\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x}) \in \mathfrak{A} \mid \mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) > 0\right\} = \mathbb{P}\Big\{\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x}) \in \mathrm{e}^{i\omega}\mathfrak{A} \;\Big|\; \mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) > 0\Big\}, \tag{A.8}$$

which yields (A.2).

Any probability measure defined on $\mathbb{S}^1$ is a Radon measure. Therefore, according to Haar's theorem [15], there exists a unique probability measure on $\mathbb{S}^1$ satisfying (A.2). Since the uniform probability measure is also invariant to phase shifts, we deduce that $\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x})$ is uniformly distributed on $\mathbb{S}^1$, conditionally to $\mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) > 0$, which concludes the proof. ■

Proposition A.2. *We assume the conditions of Proposition A.1 are met. If, moreover,* F_{X} *is strict-sense stationary, then Hypothesis 3.8 is satisfied.*

Proof. Let $n \in \mathbb{N} \setminus \{0\}$ and $\boldsymbol{x}, \boldsymbol{y}_0, \ldots, \boldsymbol{y}_{n-1} \in \mathbb{R}^2$. To alleviate notations, we consider the random vector $\boldsymbol{\mathsf{M}} = \big(\mathsf{M}_{\mathrm{X}}(\boldsymbol{y}_0), \ldots, \mathsf{M}_{\mathrm{X}}(\boldsymbol{y}_{n-1})\big)^\top$ with outcomes in $\mathbb{R}_+^n$. According to (3.53), $\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x})$ is conditionally independent of $\boldsymbol{\mathsf{M}}$ given $\mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) = 0$. Therefore, it remains to prove conditional independence given $\mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) > 0$.

The proof is organized as follows. Using a similar reasoning as Proposition A.1, we show that, for any measurable subset $\boldsymbol{\mathfrak{S}} \subset \mathbb{R}_+^n$, Z_{X} follows a uniform probability distribution conditionally to $\boldsymbol{\mathsf{M}} \in \mathfrak{S}$ and $\mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) > 0$. Since we already know that Z_{X} follows a uniform distribution conditionally to $\mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) > 0$ alone, we deduce that Z_{X} and $\boldsymbol{\mathsf{M}}$ are conditionally independent given $\mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) > 0$.

Let $\mathfrak{A} \subset \mathbb{S}^1$ and $\boldsymbol{\mathfrak{S}} := (\mathfrak{S}_i)_{i \in \{0..n-1\}} \subset \mathbb{R}_+^n$ denote measurable sets. According to (A.1), and assuming $\mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) > 0$, we get, for any $\boldsymbol{h} \in B_2(2\pi / \|\boldsymbol{\nu}\|_2)$,

$$\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x}) \in \mathfrak{A} \iff \mathcal{T}_{\boldsymbol{h}}\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x}) \in \mathrm{e}^{i\langle \boldsymbol{\nu}, \boldsymbol{h} \rangle}\mathfrak{A}; \tag{A.9}$$

$$\mathsf{M}_{\mathrm{X}}(\boldsymbol{y}_i) \in \mathfrak{S}_i \iff \mathcal{T}_{\boldsymbol{h}}\mathsf{M}_{\mathrm{X}}(\boldsymbol{y}_i) \in \mathfrak{S}_i \quad \forall i \in \{0\,..\,n-1\}\,. \tag{A.10}$$

Therefore,

$$\begin{aligned} \mathbb{P}\Big\{\big(\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x}) \in \mathfrak{A}\big) \,\&\, \big(\boldsymbol{\mathsf{M}} \in \boldsymbol{\mathfrak{S}}\big) \,\Big|\, \mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) > 0\Big\} & \\ = \mathbb{P}\left\{\big(\mathcal{T}_{\boldsymbol{h}}\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x}) \in \mathrm{e}^{i\langle \boldsymbol{\nu}, \boldsymbol{h} \rangle}\mathfrak{A}\big) \,\&\, \big(\mathcal{T}_{\boldsymbol{h}}\boldsymbol{\mathsf{M}} \in \boldsymbol{\mathfrak{S}}\big) \,\Big|\, \mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) > 0\right\}. & \end{aligned} \tag{A.11}$$

Since F_{X} is strict-sense stationary, the joint conditional probability density of

$$\mathcal{T}_{\boldsymbol{h}}\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x}),\, \mathcal{T}_{\boldsymbol{h}}\mathsf{M}_{\mathrm{X}}(\boldsymbol{y}_0),\, \ldots,\, \mathcal{T}_{\boldsymbol{h}}\mathsf{M}_{\mathrm{X}}(\boldsymbol{y}_{n-1}) \tag{A.12}$$

is identical to the one of

$$\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x}),\, \mathsf{M}_{\mathrm{X}}(\boldsymbol{y}_0),\, \ldots,\, \mathsf{M}_{\mathrm{X}}(\boldsymbol{y}_{n-1}). \tag{A.13}$$

Therefore we get

$$\begin{aligned} \mathbb{P}\Big\{\big(\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x}) \in \mathfrak{A}\big) \,\&\, \big(\boldsymbol{\mathsf{M}} \in \boldsymbol{\mathfrak{S}}\big) \,\Big|\, \mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) > 0\Big\} & \\ = \mathbb{P}\left\{\big(\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x}) \in \mathrm{e}^{i\langle \boldsymbol{\nu}, \boldsymbol{h} \rangle}\mathfrak{A}\big) \,\&\, \big(\boldsymbol{\mathsf{M}} \in \boldsymbol{\mathfrak{S}}\big) \,\Big|\, \mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) > 0\right\}. & \end{aligned} \tag{A.14}$$

We assume that $\mathbb{P}(\boldsymbol{\mathsf{M}} \in \boldsymbol{\mathfrak{S}}) > 0$. According to the above expression, and similarly to the proof of Proposition A.1, we get,

$$\begin{aligned} \forall \omega \in [0,\, 2\pi]\,,\ \mathbb{P}\Big\{\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x}) \in \mathfrak{A} \,\Big|\, \big(\boldsymbol{\mathsf{M}} \in \boldsymbol{\mathfrak{S}}\big) \,\&\, \big(\mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) > 0\big)\Big\} & \\ = \mathbb{P}\Big\{\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x}) \in \mathrm{e}^{i\omega}\mathfrak{A} \,\Big|\, \big(\boldsymbol{\mathsf{M}} \in \boldsymbol{\mathfrak{S}}\big) \,\&\, \big(\mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) > 0\big)\Big\}. & \end{aligned} \tag{A.15}$$

Then, the above conditional probability measure satisfies phase shift invariance (A.2). Therefore, as in the proof of Proposition A.1, Haar's theorem implies that $\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x})$ follows a uniform conditional distribution given $\boldsymbol{\mathsf{M}} \in \mathfrak{S}$ and $\mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) > 0$.

Moreover, strict-sense implies first-order stationarity, and thus, according to the proof of Proposition A.1, $\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x})$ follows a uniform distribution conditionally to $\mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) > 0$. Therefore we get, for any measurable sets $\mathfrak{A} \subset \mathbb{S}^1$ and $\mathfrak{S} \subset \mathbb{R}_+^n$ such that $\mathbb{P}(\boldsymbol{\mathsf{M}} \in \mathfrak{S}) > 0$,

$$\mathbb{P}\big\{\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x}) \in \mathfrak{A} \,\big|\, (\boldsymbol{\mathsf{M}} \in \mathfrak{S})\ \&\ (\mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) > 0)\big\} = \mathbb{P}\big\{\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x}) \in \mathfrak{A} \,\big|\, \mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) > 0\big\}, \tag{A.16}$$

which proves conditional independence between $\mathsf{Z}_{\mathrm{X}}(\boldsymbol{x})$ and $\boldsymbol{\mathsf{M}}$ given $\mathsf{M}_{\mathrm{X}}(\boldsymbol{x}) > 0$, and concludes the proof. ■

Remark A.3 (Stationarity hypothesis). Strict-sense stationarity suggests that any translated version of a given image is equally likely. In reality, this statement is too strong, for several reasons. First, by construction, X has all its realizations in $L^2_{\mathbb{R}}(\mathbb{R}^2)$. In that context, a stationary process yields outcomes which are zero almost everywhere. Besides, 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. More details on statistical properties of images from natural versus man-made objects can be found in a paper by Torralba and Oliva [46]. Nevertheless, this hypothesis will be considered as a reasonable approximation if the shift is much smaller than the image "characteristic" size in the continuous domain; *i.e.*, if

$$\|\boldsymbol{h}\|_2 \ll sN, \tag{A.17}$$

where, as a reminder, N denotes the support size of input images. We refer the reader to [47] for a related notion of local stationarity. As it turns out, the proofs of Propositions A.1 and A.2 only requires shifts with $\|\boldsymbol{h}\|_2 \leq 2\pi / \|\boldsymbol{\nu}\|_2$. Therefore, the constraint on $\|\boldsymbol{\theta}\|_2$ stated in (3.56) implies (A.17), and the stationarity hypothesis holds.

Remark A.4 (Justification for (A.1)). We consider

$$\mathit{\Phi}_{\mathrm{W}} : \boldsymbol{x} \mapsto \mathit{\Psi}_{\mathrm{W}}(\boldsymbol{x})\mathrm{e}^{-i\langle \boldsymbol{\nu}, \boldsymbol{x}\rangle}. \tag{A.18}$$

Similarly to Lemma 2.2, we can show that $\mathit{\Phi}_{\mathrm{W}}$ is a low-pass filter, with $\operatorname{supp} \widehat{\mathit{\Phi}}_{\mathrm{W}} \subset B_\infty(\varepsilon/2)$. For all $\boldsymbol{h} \in \mathbb{R}^2$ such that $\|\boldsymbol{h}\|_2 \leq 2\pi / \|\boldsymbol{\nu}\|_2$, we have

$$\begin{aligned}(\mathcal{T}_{\boldsymbol{h}}\mathsf{F}_{\mathrm{X}} * \overline{\mathit{\Psi}}_{\mathrm{W}})(\boldsymbol{x}) &= \iint_{\mathbb{R}^2} \mathcal{T}_{\boldsymbol{h}}\mathsf{F}_{\mathrm{X}}(\boldsymbol{x} - \boldsymbol{y})\, \overline{\mathit{\Phi}}_{\mathrm{W}}(\boldsymbol{y})\, \mathrm{e}^{-i\langle \boldsymbol{\nu}, \boldsymbol{y}\rangle}\, \mathrm{d}^2\boldsymbol{y} \\ &= \mathrm{e}^{i\langle \boldsymbol{\nu}, \boldsymbol{h}\rangle} \iint_{\mathbb{R}^2} \mathsf{F}_{\mathrm{X}}(\boldsymbol{x} - \boldsymbol{y}')\, \overline{\mathit{\Phi}}_{\mathrm{W}}(\boldsymbol{y}' - \boldsymbol{h})\, \mathrm{e}^{-i\langle \boldsymbol{\nu}, \boldsymbol{y}'\rangle}\, \mathrm{d}^2\boldsymbol{y}'.\end{aligned}$$

Since $\operatorname{supp} \widehat{\mathit{\Phi}}_{\mathrm{W}} \subset B_\infty\left(\frac{\kappa}{2s}\right)$, we can define a "minimal wavelength" $\lambda_{\mathit{\Phi}_{\mathrm{W}}} := 2\pi s/\kappa$. Then, if $\|\boldsymbol{h}\|_2 \ll \lambda_{\mathit{\Phi}_{\mathrm{W}}}$, we can approximate $\overline{\mathit{\Phi}}_{\mathrm{W}}(\boldsymbol{y}' - \boldsymbol{h}) \approx \overline{\mathit{\Phi}}_{\mathrm{W}}(\boldsymbol{y}')$. This sufficient condition is actually met, because $\|\boldsymbol{h}\|_2 \leq 2\pi / \|\boldsymbol{\nu}\|_2$ and, according to (3.57), $\|\boldsymbol{\nu}\|_2 \gg \kappa/s$. Therefore,

$$(\mathcal{T}_{\boldsymbol{h}}\mathsf{F}_{\mathrm{X}} * \overline{\mathit{\Psi}}_{\mathrm{W}})(\boldsymbol{x}) \approx \mathrm{e}^{i\langle \boldsymbol{\nu}, \boldsymbol{h}\rangle} (\mathsf{F}_{\mathrm{X}} * \overline{\mathit{\Psi}}_{\mathrm{W}})(\boldsymbol{x}). \tag{A.19}$$

As explained in Remarks A.3 and A.4, the sufficient conditions outlined in Propositions A.1 and A.2 are not strictly met. Nevertheless, we consider that Hypotheses 3.7 and 3.8 still provide a reasonable description of the distribution from which input images are drawn.

Appendix B. Details on DT-ℂWPT.

A description of the transform itself is provided in Appendix B.1. Then, Appendix B.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 B.3.

B.1. Background.

We provide a brief overview of the classical, real-valued 2D wavelet packet transform (WPT) algorithm [27, p. 377], before introducing the redundant, complex-valued and oriented DT-ℂWPT [4].

B.1.1. Discrete Wavelet Packet Transform.

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

$$\mathrm{G}_0 = \mathrm{h} \otimes \mathrm{h}; \qquad \mathrm{G}_1 = \mathrm{h} \otimes \mathrm{g}; \qquad \mathrm{G}_2 = \mathrm{g} \otimes \mathrm{h}; \qquad \mathrm{G}_3 = \mathrm{g} \otimes \mathrm{g}. \tag{B.1}$$

Let $\mathrm{X} \in l^2_{\mathbb{R}}(\mathbb{Z})$. The decomposition starts with $\mathrm{D}_0^{(0)} = \mathrm{X}$. Given $j \in \mathbb{N}$, suppose that we have computed 4^j sequences of wavelet packet coefficients at stage j, denoted by $\mathrm{D}_l^{(j)} \in l^2_{\mathbb{R}}(\mathbb{Z})$ for each $l \in \{0\,..\,4^j - 1\}$. 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 $\mathrm{D}_l^{(j)}$ into four sub-sequences, using subsampled (or strided) convolutions with kernels G_k, for each $k \in \{0\,..\,3\}$:

$$\forall k \in \{0\,..\,3\}\,,\; \mathrm{D}_{4l+k}^{(j+1)} = \left(\mathrm{D}_l^{(j)} * \overline{\mathrm{G}_k}\right) \downarrow 2. \tag{B.2}$$

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

$$\mathbf{D}^{(J)} := \left(\mathrm{D}_l^{(J)}\right)_{l\in\{0..4^J-1\}} \tag{B.3}$$

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

B.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 [18] 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 and Selesnick [4]. The latter operation, referred to as *dual-tree complex wavelet packet transform* (DT-ℂWPT), is performed as follows.

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

$$\forall \omega \in [-\pi, \pi]\,,\; \widehat{\mathrm{h}}^{[1]}(\omega) = \mathrm{e}^{-i\omega/2}\, \widehat{\mathrm{h}}^{[0]}(\omega). \tag{B.4}$$

Then, for any $k \in \{0\,..\,3\}$, we build a 2D FB $\mathbf{G}_k := (\mathrm{G}_{k,\,l})_{l \in \{0..3\}}$ similarly to (B.1):

$$\mathrm{G}_{k,\,0} = \mathrm{h}_i \otimes \mathrm{h}_j; \qquad \mathrm{G}_{k,\,1} = \mathrm{h}_i \otimes \mathrm{g}_j; \qquad \mathrm{G}_{k,\,2} = \mathrm{g}_i \otimes \mathrm{h}_j; \qquad \mathrm{G}_{k,\,3} = \mathrm{g}_i \otimes \mathrm{g}_j, \tag{B.5}$$

where $i,\, j \in \{0,\, 1\}$ are defined such that $k = 2 \times i + j$.[3]

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 $\mathbf{D}^{(J)}_{0-3}$, each of which satisfies (B.2) and (B.3). Then, for any $l \in \{0\,..\,4^J - 1\}$, the following complex feature maps are computed:

$$\begin{pmatrix} \mathrm{D}_l^{\nearrow(J)} \\ \mathrm{D}_l^{\searrow(J)} \end{pmatrix} = \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} \mathrm{D}_l^{[0](J)} \\ \mathrm{D}_l^{[3](J)} \end{pmatrix} - i \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} \mathrm{D}_l^{[2](J)} \\ \mathrm{D}_l^{[1](J)} \end{pmatrix}. \tag{B.6}$$

As explained in Appendix B.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, $\mathrm{D}_l^{\swarrow(J)}$ and $\mathrm{D}_l^{\nwarrow(J)}$ are defined as the complex conjugates of the above feature maps, and therefore do not need to be explicitly computed. Then,

$$\mathbf{D}^{(J)} := \left(\mathrm{D}_l^{\nearrow(J)},\, \mathrm{D}_l^{\searrow(J)},\, \mathrm{D}_l^{\swarrow(J)},\, \mathrm{D}_l^{\nwarrow(J)}\right)_{l \in \{0..4^J-1\}} \tag{B.7}$$

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

B.2. Convolution Operators. We now show that DT-$\mathbb{C}$WPT 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 B.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 B.1. *For any $l \in \{0\,..\,4^J - 1\}$, there exists $\mathrm{V}_l^{(J)} \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$ such that*

$$\mathrm{D}_l^{(J)} = \left(\mathrm{X} * \overline{\mathrm{V}}_l^{(J)}\right) \downarrow 2^J. \tag{B.8}$$

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

$$((\mathrm{U} \downarrow s) * \mathrm{V}) \downarrow t = (\mathrm{U} * (\mathrm{V} \uparrow s)) \downarrow (st). \tag{B.9}$$

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

$$\mathrm{V}_0^{(0)} := \mathrm{I}; \qquad \mathrm{V}_{4l+k}^{(j+1)} := \mathrm{V}_l^{(j)} * \left(\mathrm{G}_k \uparrow 2^j\right) \tag{B.10}$$

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

---

[3] Actually, the FB design requires some technicalities which are not described here.

Based on Lemma B.1, the following proposition introduces complex kernels characterizing DT-ℂWPT.

Proposition B.2. *For any* $l \in \{0\,..\,4^J-1\}$*, there exists* $\mathrm{W}_l^{\nearrow(J)} \in l^2_{\mathbb{C}}(\mathbb{Z}^2)$ *such that* (6.3) *is satisfied. Identical results are obtained with the three other Fourier quadrants.*

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

$$\mathrm{D}_l^{[m](J)} = \left(\mathrm{X} * \overline{\mathrm{V}}_l^{[m](J)}\right) \downarrow 2^J. \tag{B.11}$$

Then, the result is obtained by plugging (B.11) into (B.6) for all $m \in \{0\,..\,3\}$, and by denoting

$$\begin{pmatrix} \mathrm{W}^{\nearrow(J)}l \\ \mathrm{W}^{\searrow(J)}l \end{pmatrix} = \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} \mathrm{V}_l^{[0](J)} \\ \mathrm{V}_l^{[3](J)} \end{pmatrix} + i \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} \mathrm{V}_l^{[2](J)} \\ \mathrm{V}_l^{[1](J)} \end{pmatrix}. \tag{B.12}$$

■

Remark B.3. DT-ℂWPT, computed on a discrete image X, approximates the decomposition of a continuous 2D signal $F \in L^2_{\mathbb{R}}(\mathbb{R}^2)$ into a tight frame

$$\boldsymbol{\Psi}_{\mathbb{C}}^{(J)} := \biguplus_{l=0}^{4^J-1} \left(\Psi_{l,\boldsymbol{n}}^{\nearrow(J)},\, \Psi_{l,\boldsymbol{n}}^{\searrow(J)},\, \Psi_{l,\boldsymbol{n}}^{\swarrow(J)},\, \Psi_{l,\boldsymbol{n}}^{\nwarrow(J)}\right)_{\boldsymbol{n}\in\mathbb{Z}^2}. \tag{B.13}$$

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

$$\mathrm{D}_l^{\nearrow(J)}[\boldsymbol{n}] \approx \left(F * \overline{\Psi}_l^{\nearrow(J)*}\right)(2^J\boldsymbol{n}), \qquad \text{with} \qquad \Psi_l^{\nearrow(J)} := \Psi_{l,\boldsymbol{0}}^{\nearrow(J)}. \tag{B.14}$$

Expression (B.14) is only an approximation because of implementation technicalities that occur in practice. A “perfect” dual-tree transform should be initialized with four different inputs $\mathrm{X}^{[0-3]}$. 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 [41] for more details on this matter.

B.3. Gabor-Like Convolution Kernels. In this section, we show that the convolution kernels $\mathrm{W}_l^{\nearrow(J)}$ and $\mathrm{W}_l^{\searrow(J)}$, introduced in (6.3), approximately behave as Gabor-like filters, as defined in (2.5). To begin with, we assume that $\mathrm{h}^{[0]}$ is a Shannon filter, which is associated with a sinc scaling function [42]. 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 B.4. *There exists a permutation* $\left(\boldsymbol{\sigma}_l^{(J)}\right)_{l\in\{0..4^J-1\}}$ *of* $\{0\,..\,2^J-1\}^2$ *such that, for any* $l \in \{0\,..\,4^J-1\}$,

$$\Psi_l^{\nearrow(J)} \in \mathcal{V}\big(\boldsymbol{\theta}_l^{(J)},\, \kappa_J\big), \tag{B.15}$$

where $\Psi_l^{\nearrow(J)}$ *has been introduced in Remark* B.3*, and where we have defined*

$$\boldsymbol{\theta}_l^{(J)} := \left(\boldsymbol{\sigma}_l^{(J)} + \frac{1}{2}\right)\frac{\pi}{2^J} \qquad \textit{and} \qquad \kappa_J := \frac{\pi}{2^J}. \tag{B.16}$$

We remind the reader that $\mathcal{V}\big(\boldsymbol{\nu},\, \varepsilon\big)$*, defined in* (2.2)*, denotes a space of Gabor-like filters in the continuous framework.*

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

$$\Psi_l^{\nearrow(J)} = \psi_{l_1}^{(J)} \otimes \psi_{l_2}^{(J)}, \tag{B.17}$$

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

$$\psi_{l'}^{(J)} = \psi_{l'}^{[0](J)} + i\,\psi_{l'}^{[1](J)}, \tag{B.18}$$

where $\psi_{l'}^{[0](J)} \in L^2_{\mathbb{R}}(\mathbb{R})$ 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

$$\widehat{\psi}_{l'}^{(J)} = 2\,\widehat{\psi}_{l'}^{[0](J)}\,\mathbb{1}_{\mathbb{R}_+}. \tag{B.19}$$

Consequently, according to the Coifman-Wickerhauser theorem [27, pp. 384-385], there exists $k \in \big\{0\,..\,2^J - 1\big\}$ such that

$$\operatorname{supp} \widehat{\psi}_{l'}^{(J)} \subset \left[\frac{k\pi}{2^J},\, \frac{(k+1)\pi}{2^J}\right]. \tag{B.20}$$

Finally, the tensor product (B.17) yields the result. ■

According to Proposition B.4, each atom $\Psi_l^{\nearrow(J)}$, for $l \in \big\{0\,..\,4^J - 1\big\}$, 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 B.4 that the discrete filter $\mathrm{W}_l^{\nearrow(J)} \in l^2_{\mathbb{C}}(\mathbb{Z}^2)$ satisfies the Gabor property (6.4). However, as mentioned in Remark B.3, (B.14) is only an approximation. In fact, the Fourier support of $\mathrm{W}_l^{\nearrow(J)}$ 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 [41] yields near-analytic solutions even for small values of J. We therefore consider (6.4) as a reasonable approximation if $J \geq 2$.

Remark B.5. Proposition B.4 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 $\mathrm{W}_l^{\nearrow(J)}$ 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 [30], or finite-length filters that approximate the half-sample delay condition [41]. Therefore, energy is leaking outside the square cells tiling the Fourier domain. To counterbalance this, we increase the window size up to

$$\kappa_J := \frac{\pi}{2^{J-1}} = \pi/m_J, \tag{B.21}$$

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 $\kappa_J := \pi/2^{J-1}$.

and consider that (6.4) remains a reasonable approximation. Therefore, the conditions to apply Theorems 2.9, 3.11, and 4.2 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:

$$\rho_l^{\nearrow} := \frac{\max_{\boldsymbol{\theta}\in[-\pi,\,\pi]^2} \left\| \mathbb{1}_{B_\infty(\boldsymbol{\theta},\,\kappa_J/2)} \widehat{\mathrm{W}}_l^{\nearrow(J)} \right\|_{L^2}^2}{\left\| \widehat{\mathrm{W}}_l^{\nearrow(J)} \right\|_{L^2}^2}, \tag{B.22}$$

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 2.1. If (6.4) is perfectly satisfied, then $\rho_l^{\nearrow} = 1$. The statistics computed over the collection $\left(\rho_l^{\nearrow},\, \rho_l^{\searrow}\right)_{l\in\{0..4^J-1\}}$ are reported in Table 1.

Remark B.6. For "boundary filters", *i.e.*, when $\left\|\boldsymbol{\theta}_l^{(J)}\right\|_\infty = \left(1 - 2^{-(J+1)}\right)\pi$, Remark 2.1 states that a small fraction of the filter's energy remains located at the far end of the Fourier domain—see also [4]. Therefore, these filters do not strictly comply with the conditions of Theorems 2.9, 3.11, and 4.2. We nevertheless include them in our experiments.

REFERENCES

[1] K. B. Athreya and S. N. Lahiri, *Measure Theory and Probability Theory*, vol. 19, Springer, 2006.

[2] A. Azulay and Y. Weiss, *Why do deep convolutional networks generalize so poorly to small image transformations?*, Journal of Machine Learning Research, 20 (2019), pp. 1–25.

[3] R. Balan, M. Singh, and D. Zou, *Lipschitz Properties for Deep Convolutional Networks*, Contemporary Mathematics, 706 (2018), pp. 129–151.

[4] I. Bayram and I. W. Selesnick, *On the Dual-Tree Complex Wavelet Packet and M-Band Transforms*, IEEE Transactions on Signal Processing, 56 (2008), pp. 2298–2310.

[5] A. Bietti, *Approximation and Learning with Deep Convolutional Models: A Kernel Perspective*, in International Conference on Learning Representations, 2022.

[6] A. Bietti and J. Mairal, *Invariance and Stability of Deep Convolutional Representations*, in Advances in Neural Information Processing Systems, 2017, p. 1622.

[7] A. Bietti and J. Mairal, *Group invariance, stability to deformations, and complexity of deep convolutional representations*, The Journal of Machine Learning Research, 20 (2019), pp. 876–924.

[8] A. Bietti and J. Mairal, *On the Inductive Bias of Neural Tangent Kernels*, in Advances in Neural Information Processing Systems, vol. 32, Curran Associates, Inc., 2019.

[9] J. Bruna and S. Mallat, *Invariant Scattering Convolution Networks*, IEEE Transactions on Pattern Analysis and Machine Intelligence, 35 (2013), pp. 1872–1886.

[10] F. Cotter and N. Kingsbury, *A Learnable Scatternet: Locally Invariant Convolutional Layers*, in 2019 IEEE International Conference on Image Processing (ICIP), 2019, pp. 350–354.

[11] W. Czaja and W. Li, *Analysis of time-frequency scattering transforms*, Applied and Computational Harmonic Analysis, 47 (2019), pp. 149–171.

[12] W. Czaja and W. Li, *Rotationally Invariant Time–Frequency Scattering Transforms*, Journal of Fourier Analysis and Applications, 26 (2020), p. 4.

[13] F. Gama, A. Ribeiro, and J. Bruna, *Stability of Graph Scattering Transforms*, in Advances in Neural Information Processing Systems, vol. 32, Curran Associates, Inc., 2019.

[14] S. Gauthier, B. Thérien, L. Alsène-Racicot, M. Chaudhary, I. Rish, E. Belilovsky, M. Eickenberg, and G. Wolf, *Parametric Scattering Networks*, in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 5749–5758.

[15] P. R. Halmos, *Measure Theory*, Springer, 2013.

[16] J. Havlicek, J. Havlicek, and A. Bovik, *The analytic image*, in Proceedings of International Conference on Image Processing, vol. 2, 1997, pp. 446–449 vol.2.

[17] K. Huang and S. Aviyente, *Wavelet feature selection for image classification*, IEEE Transactions on Image Processing, 17 (2008), pp. 1709–1720.

[18] N. Kingsbury, *Complex wavelets for shift invariant analysis and filtering of signals*, Applied and computational harmonic analysis, 10 (2001), pp. 234–253.

[19] N. Kingsbury, *Design of Q-shift complex wavelets for image processing using frequency domain energy minimization*, in Proceedings International Conference on Image Processing, vol. 1, 2003, pp. I–1013.

[20] N. Kingsbury and J. Magarey, *Wavelet Transforms in Image Processing*, in Signal Analysis and Prediction, Applied and Numerical Harmonic Analysis, Birkhäuser, Boston, MA, 1998, pp. 27–46.

[21] A. Laine and J. Fan, *Texture Classification by Wavelet Packet Signatures*, IEEE Transactions on Pattern Analysis and Machine Intelligence, 15 (1993), pp. 1186–1191.

[22] Y. LeCun, Y. Bengio, and G. Hinton, *Deep learning*, Nature, 521 (2015), pp. 436–444.

[23] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, *Gradient-based learning applied to document recognition*, Proceedings of the IEEE, 86 (1998), pp. 2278–2323.

[24] H. Leterme, K. Polisano, V. Perrier, and K. Alahari, *Modélisation Parcimonieuse de CNNs avec des Paquets d'Ondelettes Dual-Tree*, in ORASIS, 2021.

[25] H. Leterme, K. Polisano, V. Perrier, and K. Alahari, *From CNNs to Shift-Invariant Twin Models Based on Complex Wavelets*, 2023, https://arxiv.org/abs/2212.00394.

[26] B. Liao and F. Peng, *Rotation-invariant texture features extraction using Dual-Tree Complex Wavelet Transform*, in 2010 International Conference on Information, Networking and Automation (ICINA), vol. 1, 2010.

[27] S. Mallat, *A Wavelet Tour of Signal Processing : The Sparse Way*, Academic Press, 2009.

[28] S. Mallat, *Group invariant scattering*, Communications on Pure and Applied Mathematics, 65 (2012), pp. 1331–1398.

[29] S. Mallat, *Understanding deep convolutional networks*, Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 374 (2016), p. 20150203.

[30] Y. Meyer, *Principe d'incertitude, bases hilbertiennes et algèbres d'opérateurs*, in Séminaire Bourbaki, vol. 662, Société Mathematique de France, 1985.

[31] G. P. Nason and B. W. Silverman, *The Stationary Wavelet Transform and some Statistical Applications*, in Wavelets and Statistics, Lecture Notes in Statistics, Springer, 1995, pp. 281–299.

[32] E. Oyallon, *Analyzing and Introducing Structures in Deep Convolutional Neural Networks*, doctoral thesis, Paris Sciences et Lettres, 2017.

[33] E. Oyallon, E. Belilovsky, S. Zagoruyko, and M. Valko, *Compressing the Input for CNNs with the First-Order Scattering Transform*, in Proceedings of the European Conference on Computer Vision (ECCV), 2018.

[34] K. I. Park and M. Park, *Fundamentals of Probability and Stochastic Processes with Applications to Communications*, Springer, 2018.

[35] J. C. Pérez, M. Alfarra, G. Jeanneret, A. Bibi, A. Thabet, B. Ghanem, and P. Arbeláez, *Gabor Layers Enhance Network Robustness*, in Proceedings of the European Conference on Computer Vision (ECCV), 2020.

[36] M. Perlmutter, F. Gao, G. Wolf, and M. Hirn, *Geometric Wavelet Scattering Networks on Compact Riemannian Manifolds*, in Proceedings of The First Mathematical and Scientific Machine Learning Conference, PMLR, 2020, pp. 570–604.

[37] S. Pittner and S. V. Kamarthi, *Feature extraction from wavelet coefficients for pattern recognition tasks*, IEEE Transactions on Pattern Analysis and Machine Intelligence, 21 (1999), pp. 83–88.

[38] M. Rai and P. Rivas, *A Review of Convolutional Neural Networks and Gabor Filters in Object Recognition*, in 2020 International Conference on Computational Science and Computational Intelligence (CSCI), 2020, pp. 1560–1567.

[39] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, *ImageNet Large Scale Visual Recognition Challenge*, International Journal of Computer Vision, 115 (2015), pp. 211–252.

[40] M. Scetbon and Z. Harchaoui, *Harmonic Decompositions of Convolutional Networks*, in International Conference on Machine Learning, PMLR, 2020, pp. 8522–8532.

[41] I. W. Selesnick, R. Baraniuk, and N. Kingsbury, *The dual-tree complex wavelet transform*, IEEE Signal Processing Magazine, 22 (2005), pp. 123–151.

[42] C. Shannon, *Communication in the Presence of Noise*, Proceedings of the IRE, 37 (1949), pp. 10–21.

[43] L. Sifre and S. Mallat, *Rotation, Scaling and Deformation Invariant Scattering for Texture Discrimination*, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2013, pp. 1233–1240.

[44] A. Singh and N. 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.

[45] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, *Intriguing properties of neural networks*, 2014, https://arxiv.org/abs/1312.6199.

[46] A. Torralba and A. Oliva, *Statistics of natural image categories*, Network: Computation in Neural Systems, 14 (2003), pp. 391–412.

[47] M. Tygert, J. Bruna, S. Chintala, Y. LeCun, S. Piantino, and A. Szlam, *A Mathematical Motivation for Complex-Valued Convolutional Networks*, Neural Computation, 28 (2016), pp. 815–825.

[48] M. Vetterli, *Wavelets, approximation, and compression*, IEEE Signal Processing Magazine, 18 (2001), pp. 59–73.

[49] A. Virmaux and K. Scaman, *Lipschitz regularity of deep neural networks: Analysis and efficient estimation*, in Advances in Neural Information Processing Systems, 2018.

[50] I. Waldspurger, *Wavelet Transform Modulus : Phase Retrieval and Scattering*, doctoral thesis, Ecole normale supérieure, Paris, 2015.

[51] T. Wiatowski and H. Bölcskei, *A Mathematical Theory of Deep Convolutional Neural Networks for Feature Extraction*, IEEE Transactions on Information Theory, 64 (2018), pp. 1845–1866.

[52] G. G. Yen, *Wavelet packet feature extraction for vibration monitoring*, IEEE Transactions on Industrial Electronics, 47 (2000), pp. 650–667.

[53] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, *How transferable are features in deep neural networks?*, in Advances in Neural Information Processing Systems, 2014, pp. 3320–3328.

[54] J. Zarka, F. Guth, and S. Mallat, *Separation and Concentration in Deep Networks*, in International Conference on Learning Representations, 2021.

[55] J. Zarka, L. Thiry, T. Angles, and S. Mallat, *Deep Network Classification by Scattering and Homotopy Dictionary Learning*, in International Conference on Learning Representations, 2020.

[56] R. Zhang, *Making Convolutional Networks Shift-Invariant Again*, in International Conference on Machine Learning, 2019.

[57] D. Zou, R. Balan, and M. Singh, *On Lipschitz Bounds of General Convolutional Neural Networks*, IEEE Transactions on Information Theory, 66 (2020), pp. 1738–1759.

[58] D. Zou and G. Lerman, *Graph convolutional neural networks via scattering*, Applied and Computational Harmonic Analysis, 49 (2020), pp. 1046–1074.

[59] X. Zou, F. Xiao, Z. Yu, Y. Li, and Y. J. Lee, *Delving Deeper into Anti-Aliasing in ConvNets*, International Journal of Computer Vision, 131 (2023), pp. 67–81.