

Next, we show that, after training with ImageNet, the feature extractor formed by early layers of popular CNN architectures can approximately be reformulated as a stack of \mathbb{R} GPool operators. Our framework therefore provides a theoretical grounding to study these networks.

We apply our theoretical results on the dual-tree complex wavelet packet transform (DT-CWPT), a particular case of discrete Gabor-like decomposition with perfect reconstruction properties [29], [30], possessing characteristics comparable to standard convolution layers. Finally, we verify our predictions on a deterministic setting based on DT-CWPT. Given an input image, we compute the mean discrepancy between the outputs of CGMod and \mathbb{R} GPool, for each wavelet packet filter.¹ We then observe that shift invariance, when measured on \mathbb{R} GPool feature maps, is nearly achieved if they remain close to CGMod outputs. We therefore establish an invariance validity domain for \mathbb{R} GPool operators.

Prior to this work, we presented a preliminary study [31], where we experimentally showed that an operator based on 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,² which are known to be generally unstable [32]. Yet, we observed a limited but genuine form of shift invariance, compared to other models based on the standard, non-analytic wavelet packet transform. At the same time, we became aware of a preliminary work in Waldspurger’s PhD thesis [28, 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 discrete and subsampled convolutions. We address this issue by adopting a probabilistic point of view.

II. RELATED WORK

A. Wavelet Scattering Networks

These models, introduced by Bruna and Mallat [11], compute cascading wavelet convolutions followed by non-linear operations. They produce translation-invariant image representations which are stable to deformation and preserve high-frequency information. A variation has been proposed in [33] to improve stability with respect to small rotations. Wavelet scattering networks were later adapted to the discrete framework using the dual-tree complex wavelet transform [34], as well as functions defined on graphs [35].

Such networks, which are totally deterministic aside from the output classifier, achieve results on small image datasets but do not scale well to more complex ones. According to Oyallon *et al.* [36], [37], this is partly due to non-geometric sources of variability within classes. Instead, the authors proposed to use scattering coefficients as inputs to a CNN,

showing that the network complexity can be reduced while keeping competitive performance. More recent work by Zarka *et al.* [38] proposed to sparsify wavelet scattering coefficients by learning a dictionary matrix, and managed to outperform AlexNet [10]. This was extended by the same team in [39], where the authors proposed to learn 1×1 convolutions between feature maps of scattering coefficients and to apply soft-thresholding to reduce within-class variability. This model reached the classification accuracy of ResNet-18 on ImageNet.

Other work proposed architectures in which the scattering transform is no longer deterministic. Cotter and Kingsbury [40] built a learnable scattering network. In this model, feature maps of scattering coefficients are mixed together using trainable weights, to account for cross-channel filtering as implemented in CNNs. Their architecture outperformed VGG networks on small image datasets. Recently, Gauthier *et al.* [41] introduced parametric scattering networks, in which the scale, orientation and aspect ratio of each wavelet filter are adjusted during training. Their approach has proven successful when trained on limited dataset.

All these papers are driven by the purpose of building ad-hoc CNN-like feature extractors, implementing well defined mathematical operators specifically designed to meet a certain number of desired properties. By contrast, our work seeks evidence that such properties, which have been established for wavelet scattering networks, are—to some extent—embedded in *existing* CNN architectures, with no need to alter their behavior or introduce new features.

B. Invariance Studies in CNNs

Several papers analyze invariance properties in CNN-related feature extractors, including—but not limited to—wavelet scattering networks. Whereas extensive studies related to the original architecture are proposed by Mallat in [42], [43], more recent work tackle the question for various extensions of the model. In [44], [45], scattering networks based on uniform covering frames—i.e., frames splitting the frequency domain into windows of roughly equal size, much like Gabor frames—are studied. Besides, [8] considers a wide variety of feature extractors involving convolutions, Lipschitz-continuous nonlinearities and pooling operators. The paper shows that outputs become more translation invariant with increasing network depth. Finally, [46] shows 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.

In these studies, invariance properties are obtained for continuous signals. Whereas real-life CNNs can be mathematically described in the continuous framework, feature maps computed at their hidden and output layers are actually discrete sequences, which can be recovered by sampling the continuous signals. At each convolution and pooling layers, the sampling interval is increased (subsampling), resulting in a loss of information. Unfortunately, this may greatly affect shift invariance, as explained in section I. The current paper specifically addresses this issue.

¹DT-CWPT paves the Fourier domain into square regions of identical size, each of them associated to a specific filter.

²To do so, we split the real and imaginary parts of the original filters.

III. SHIFT INVARIANCE OF OPERATORS

The goal of this section is to theoretically establish conditions for near-shift invariance at the output of the first max pooling layer. We start by proving shift invariance of $\mathbb{C}\text{GMod}$ operators. Then, we establish conditions under which $\mathbb{R}\text{GPool}$ and $\mathbb{C}\text{GMod}$ produce closely related outputs. Finally, we derive a probabilistic measure of shift invariance for $\mathbb{R}\text{GPool}$.

A. Notations

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

Continuous Framework: Given $p > 0$ and a measurable subset of \mathbb{R} or \mathbb{R}^2 denoted by E , we consider $L^p(E)$ as the space of measurable complex-valued functions $f : E \rightarrow \mathbb{C}$ such that $\|f\|_{L^p} := \int_E |f(x)|^p dx < +\infty$. Whenever we talk about equality in $L^p(E)$ or inclusion in E , it shall be understood as ‘‘almost everywhere with respect to the Lebesgue measure’’. Besides, we denote by $L_{\mathbb{R}}^2(\mathbb{R}^2) \subset L^2(\mathbb{R}^2)$ the subset of real-valued functions. For any $f \in L^2(\mathbb{R}^2)$, \bar{f} denotes its flipped version: $\bar{f}(x) := f(-x)$.

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

$$\forall \nu \in \mathbb{R}^2, \hat{f}(\nu) := \iint_{\mathbb{R}^2} f(x) e^{-i\langle \nu, x \rangle} d^2x. \quad (1)$$

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

$$\mathcal{V}(\nu, \varepsilon) := \left\{ \psi \in L^2(\mathbb{R}^2) \mid \text{supp } \hat{\psi} \subset \mathcal{B}_{\infty}(\nu, \varepsilon/2) \right\}. \quad (2)$$

For any $\mathbf{h} \in \mathbb{R}^2$, we also consider the translation operator, denoted by $\mathcal{T}_{\mathbf{h}}$, defined by $\mathcal{T}_{\mathbf{h}}f : x \mapsto f(x - \mathbf{h})$.

Discrete Framework: We consider $l^2(\mathbb{Z}^d)$ as the space of d -dimensional sequences $X \in \mathbb{C}^{\mathbb{Z}^d}$ such that $\|X\|_2^2 := \sum_{\mathbf{n} \in \mathbb{Z}^d} |X[\mathbf{n}]|^2 < +\infty$. Indexing is made between square brackets: $\forall X \in l^2(\mathbb{Z}^d), \forall \mathbf{n} \in \mathbb{Z}^d, X[\mathbf{n}] \in \mathbb{C}$, and we denote by $l_{\mathbb{R}}^2(\mathbb{Z}^d) \subset l^2(\mathbb{Z}^d)$ the subset of real-valued sequences. For any $X \in l^2(\mathbb{Z}^d)$, \bar{X} denotes its flipped version: $\bar{X}[\mathbf{n}] := X[-\mathbf{n}]$. The subsampling operator is denoted by \downarrow : for any $X \in l^2(\mathbb{Z}^d)$ and any $m \in \mathbb{N}^*$, $(X \downarrow m)[\mathbf{n}] := X[m\mathbf{n}]$.

2D images, feature maps and convolution kernels are considered as elements of $l^2(\mathbb{Z}^2)$, and are denoted by straight capital letters. Besides, arrays of 2D sequences are denoted by bold straight capital letters, for instance: $\mathbf{X} = (X_k)_{k \in \{0..K-1\}}$. Note that indexing starts at 0 to comply with practical implementations. We will also consider 1D sequences $x \in l^2(\mathbb{Z})$, denoted by straight lower-case letters.

The 2D discrete-time Fourier transform of any $X \in l^2(\mathbb{Z}^2)$ is denoted by $\hat{X} \in L^2([-\pi, \pi]^2)$, such that

$$\forall \xi \in [-\pi, \pi]^2, \hat{X}(\xi) := \sum_{\mathbf{n} \in \mathbb{Z}^2} X[\mathbf{n}] e^{-i\langle \xi, \mathbf{n} \rangle}. \quad (3)$$




Fig. 1. Spatial (left) and Fourier (right) representations of convolution kernels in the first layer of AlexNet, after training with ImageNet ILSVRC2012. Each kernel connects the 3 RGB input channels to one of the 64 output channels.

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

$$\mathcal{G}(\xi, \kappa) := \left\{ W \in l^2(\mathbb{Z}^2) \mid \text{supp } \widehat{W} \subset \mathcal{B}_{\infty}(\xi, \kappa/2) \right\}. \quad (4)$$

Remark 1: The support $\mathcal{B}_{\infty}(\xi, \kappa/2)$ actually lives in the quotient space $[-\pi, \pi]^2 / (2\pi\mathbb{Z}^2)$. Consequently, when ξ 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 ξ and κ is implicitly assumed to avoid such a situation.

B. 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. Fig. 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. We denote by $V \in l_{\mathbb{R}}^2(\mathbb{Z}^2)$ one of these ‘‘well-behaved’’ filters. Its Fourier spectrum roughly consists in two bright spots which are symmetric with respect to the origin.³ Now, we consider a complex-valued companion $W \in l^2(\mathbb{Z}^2)$ such that, for any $\xi = (\xi_1, \xi_2) \in [-\pi, \pi]^2$,

$$\widehat{W}(\xi) := (1 + \text{sgn } \xi_1) \cdot \widehat{V}(\xi). \quad (5)$$

We can show that V is the real part of W , and that $W = V + i\mathcal{H}(V)$, where \mathcal{H} denotes the two-dimensional Hilbert transform as introduced in [47].⁴ As a consequence, \widehat{W} is equal to $2\widehat{V}$ on one half of the Fourier domain, and 0 on the other half. Therefore, only one bright spot remains in the spectrum. It turns out that such complex filters with high frequency resolution produce stable signal representations, as we will see in section III-C. In the subsequent sections, we then wonder whether this property is kept when considering the max pooling of real-valued convolutions.

³Actually, the Fourier transform of any real-valued sequence is centrally symmetric: $\widehat{V}(-\xi) = \widehat{V}(\xi)^*$. The specificity of well-oriented filters lies in the concentration of their power spectrum around two precise locations.

⁴ $\mathcal{H}(V)$ is defined such that $\mathcal{H}(\widehat{V})(\xi) := -i \text{sgn}(\xi_1) \widehat{V}(\xi)$.

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.

C. Shift Invariance of CGMod Outputs

The aim of this section is to show that the modulus of discrete Gabor-like coefficients—i.e., the output of a CGMod operator such as introduced in section I-B—is nearly shift-invariant (the meaning of shift invariance will be clarified). This result is hinted in [32] but not formally proven.

1) *Continuous Framework:* We introduce several results regarding functions defined on the continuous space \mathbb{R}^2 . Near-shift invariance on discrete 2D sequences will then be derived from these results by taking advantage of sampling theorems. Lemma 1 below is adapted from [28, pp. 190–191].

Lemma 1: Given $\varepsilon > 0$ and $\nu \in \mathbb{R}^2$, let $\psi \in \mathcal{V}(\nu, \varepsilon)$ denote a complex-valued filter such as defined in (2). Now, for any real-valued function $f \in L^2_{\mathbb{R}}(\mathbb{R}^2)$, we consider the complex-valued function $f_0 \in L^2(\mathbb{R}^2)$ defined by

$$f_0 : \mathbf{x} \mapsto (f * \bar{\psi})(\mathbf{x}) e^{i\langle \nu, \mathbf{x} \rangle}. \quad (6)$$

Then f_0 is low-frequency, with $\text{supp } \widehat{f_0} \subset \mathcal{B}_{\infty}(\varepsilon/2)$.

Proof: See Appendix A. ■

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

Proposition 1: For any $f_0 \in L^2_{\mathbb{R}}(\mathbb{R}^2)$ such that $\text{supp } \widehat{f_0} \subset \mathcal{B}_{\infty}(\varepsilon/2)$, and any $\mathbf{h} \in \mathbb{R}^2$ satisfying $\|\mathbf{h}\|_1 \leq \pi/\varepsilon$,

$$\|\mathcal{T}_{\mathbf{h}} f_0 - f_0\|_{L^2} \leq \alpha(\varepsilon \mathbf{h}) \|f_0\|_{L^2}, \quad (7)$$

where we have defined

$$\alpha : \boldsymbol{\tau} \mapsto \frac{\|\boldsymbol{\tau}\|_1}{2}. \quad (8)$$

Proof: See Appendix B. ■

2) *Adaptation to Discrete 2D Sequences:* Given $\kappa \in]0, 2\pi]$ and $\boldsymbol{\xi} \in \mathcal{B}_{\infty}(\pi)$, let $W \in \mathcal{G}(\boldsymbol{\xi}, \kappa)$ denote a discrete Gabor-like filter such as defined in (4). For any image $X \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$ with finite support and any subsampling factor $m \in \mathbb{N}^*$, we express $(X * \bar{W}) \downarrow m$ using the continuous framework introduced above, and derive an invariance formula.

For any sampling interval $s \in \mathbb{R}^*_+$, let \mathcal{U}_s denote the space of 2D functions $g \in L^2(\mathbb{R}^2)$ such that the support of \widehat{g} is included in $\mathcal{B}_{\infty}(\pi/s)$.⁵ We consider the following lemma.

Lemma 2: Let $s > 0$. For any $g \in \mathcal{U}_s$ and any $\boldsymbol{\omega} \in \mathcal{B}_{\infty}(\pi/s)$, we have

$$\widehat{g}(\boldsymbol{\omega}) = s \widehat{Y}(s\boldsymbol{\omega}), \quad (9)$$

where $Y \in l^2(\mathbb{Z}^2)$ is defined such that $Y[\mathbf{n}] := s g(s\mathbf{n})$, for any $\mathbf{n} \in \mathbb{Z}^2$. Besides, we have the following norm equality:

$$\|g\|_{L^2} = \|Y\|_2. \quad (10)$$

Proof: See Appendix C. ■

We now consider $\phi^{(s)} \in L^2_{\mathbb{R}}(\mathbb{R}^2)$ such that $\widehat{\phi^{(s)}} := s \mathbb{1}_{\mathcal{B}_{\infty}(\pi/s)}$.⁶ For any $\mathbf{n} \in \mathbb{Z}^2$, we denote by $\phi_{\mathbf{n}}^{(s)} := \mathcal{T}_{s\mathbf{n}} \phi^{(s)}$ a shifted version of $\phi^{(s)}$. According to Theorem 3.5 in [48, p. 68], $\{\phi_{\mathbf{n}}^{(s)}\}_{\mathbf{n} \in \mathbb{Z}^2}$ is an orthonormal basis of \mathcal{U}_s .⁷ We then get the following proposition, which draws a bond between the discrete and continuous frameworks.

Proposition 2: Let $X \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$ denote an input image with finite support, and $W \in \mathcal{G}(\boldsymbol{\xi}, \kappa)$. Considering a sampling interval $s \in \mathbb{R}^*_+$, we define $f_X \in L^2_{\mathbb{R}}(\mathbb{R}^2)$ and $\psi_W \in L^2(\mathbb{R}^2)$ such that

$$f_X := \sum_{\mathbf{n} \in \mathbb{Z}^2} X[\mathbf{n}] \phi_{\mathbf{n}}^{(s)} \quad \text{and} \quad \psi_W := \sum_{\mathbf{n} \in \mathbb{Z}^2} W[\mathbf{n}] \phi_{\mathbf{n}}^{(s)}. \quad (11)$$

Then, $\psi_W \in \mathcal{V}(\boldsymbol{\xi}/s, \kappa/s)$. Moreover, for all $\mathbf{n} \in \mathbb{Z}$,

$$X[\mathbf{n}] = s f_X(s\mathbf{n}); \quad W[\mathbf{n}] = s \psi_W(s\mathbf{n}), \quad (12)$$

and, for a given subsampling factor $m \in \mathbb{N}^*$,

$$\left((X * \bar{W}) \downarrow m \right) [\mathbf{n}] = (f_X * \bar{\psi}_W)(m s \mathbf{n}). \quad (13)$$

Proof: See Appendix D. ■

Proposition 2 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 $\mathbf{u} \in \mathbb{R}^2$, a translation operator $\mathcal{T}_{\mathbf{u}}$ on discrete sequences, even if \mathbf{u} has non-integer values:

$$\mathcal{T}_{\mathbf{u}} X[\mathbf{n}] := s \mathcal{T}_{s\mathbf{u}} f_X(s\mathbf{n}), \quad (14)$$

where f_X is defined in (11). We can indeed show that this definition is independent from the choice of sampling interval $s > 0$. Besides, given $X \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$, we have

$$\forall \mathbf{k} \in \mathbb{Z}^2, \mathcal{T}_{\mathbf{k}} X[\mathbf{n}] = X[\mathbf{n} - \mathbf{k}]; \quad (15)$$

$$\forall \mathbf{u}, \mathbf{v} \in \mathbb{R}^2, \mathcal{T}_{\mathbf{u}}(\mathcal{T}_{\mathbf{v}} X) = \mathcal{T}_{\mathbf{u}+\mathbf{v}} X, \quad (16)$$

which shows that $\mathcal{T}_{\mathbf{u}}$ corresponds to the intuitive idea of a shift operator. Expressions (15) and (16) are direct consequence of the following lemma, which bonds the shift operator in the discrete and continuous frameworks.

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

$$f_{\mathcal{T}_{\mathbf{u}} X} = \mathcal{T}_{s\mathbf{u}} f_X. \quad (17)$$

Proof: See Appendix E. ■

We now consider the following corollary to Proposition 2.

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

$$\left((\mathcal{T}_{\mathbf{u}} X * \bar{W}) \downarrow m \right) [\mathbf{n}] = (\mathcal{T}_{s\mathbf{u}} f_X * \bar{\psi}_W)(m s \mathbf{n}). \quad (18)$$

Proof: Apply (13) in Proposition 2 with $X \leftarrow \mathcal{T}_{\mathbf{u}} X$, and use Lemma 3 to conclude. ■

⁶ $\phi^{(s)}$ is a tensor product of scaled and normalized sinc functions.

⁷We actually use the 2D formulation, mentioned in p. 82.

⁵Using the notation introduced in (2), we have $\mathcal{U}_s = \mathcal{V}(\mathbf{0}, 2\pi/s)$.

3) *Shift Invariance in the Discrete Framework:* We consider the following operator, for any $W \in l^2(\mathbb{Z}^2)$:

$$\mathcal{C}_m[W] : X \mapsto |(X * \overline{W}) \downarrow m|. \quad (19)$$

When $W \in \mathcal{G}(\xi, \kappa)$, we refer to this as a $\mathbb{C}\text{GMod}$ operator. For the sake of concision, in what follows we will write \mathcal{C}_m instead of $\mathcal{C}_m[W]$, when no ambiguity is possible.

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

Theorem 1: Let $W \in \mathcal{G}(\xi, \kappa)$ denote a discrete Gabor-like filter and $m \in \mathbb{N}^*$ denote a subsampling factor. If $\kappa \leq 2\pi/m$, then, for any input image $X \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$ with finite support and any translation vector $\mathbf{u} \in \mathbb{R}^2$ satisfying $\|\mathbf{u}\|_1 \leq \pi/\kappa$,

$$\|\mathcal{C}_m(\mathcal{T}_{\mathbf{u}}X) - \mathcal{C}_mX\|_2 \leq \alpha(\kappa\mathbf{u}) \|\mathcal{C}_mX\|_2, \quad (20)$$

where α has been defined in (8).

Proof: The proof of this theorem, which involves Lemmas 1-2, Propositions 1-2 and Corollary 1, is provided in Appendix F. ■

Interestingly, the reference value used in Theorem 1, i.e., $\|\mathcal{C}_mX\|_2$, is fully shift-invariant, as stated in the following proposition.

Proposition 3: Let $W \in \mathcal{G}(\xi, \kappa)$ and $m \in \mathbb{N}^*$. Assuming $\kappa \leq 2\pi/m$, we have, for any $X \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$ and any $\mathbf{u} \in \mathbb{R}^2$,

$$\|\mathcal{C}_m(\mathcal{T}_{\mathbf{u}}X)\|_2 = \|\mathcal{C}_mX\|_2. \quad (21)$$

Proof: See Appendix G. ■

D. From $\mathbb{C}\text{GMod}$ to $\mathbb{R}\text{GPool}$

Since $\mathbb{C}\text{GMod}$ operators are not found in classical CNN architectures, the above result does not apply straightforwardly. Instead, the first convolution layer contains real-valued kernels, and is generally followed by ReLU and max pooling. As shown in section IV, this process can be described as an operator parameterized by $W \in l^2(\mathbb{Z}^2)$, defined by

$$\mathcal{R}_{m,q}[W] : X \mapsto \text{MaxPool}_q \left((X * \overline{\text{Re}W}) \downarrow m \right), \quad (22)$$

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

$$\text{MaxPool}_q(Y)[\mathbf{n}] := \max_{\|\mathbf{k}\|_{\infty} \leq q} Y[2\mathbf{n} + \mathbf{k}]. \quad (23)$$

As hinted in section III-B, an important number of trained convolution kernels exhibit oscillating patterns with various scales and orientations. In such a case, $W \in \mathcal{G}(\xi, \kappa)$ for a certain value of $\xi \in [-\pi, \pi]^2$ and $\kappa \in]0, 2\pi]$, and we refer to $\mathcal{R}_{m,q}[W]$ as an $\mathbb{R}\text{GPool}$ operator. For the sake of concision, from now on we write $\mathcal{R}_{m,q}$ instead of $\mathcal{R}_{m,q}[W]$, when no ambiguity is possible.

In what follows, we show that, under specific conditions on \widehat{W} , $\mathbb{R}\text{GPool}$ and $\mathbb{C}\text{GMod}$ operators produce comparable outputs. We then provide a shift invariance bound for $\mathbb{R}\text{GPool}$.

1) *Continuous Framework:* This paragraph, directly adapted from [28, pp. 190–191], provides an intuition about resemblance between $\mathbb{R}\text{GPool}$ and $\mathbb{C}\text{GMod}$ in the continuous framework. As will be highlighted later in this section III-D, 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 $\psi \in \mathcal{V}(\nu, \varepsilon)$. Let us also consider

$$g : (\mathbf{x}, \mathbf{h}) \mapsto \cos(\langle \nu, \mathbf{h} \rangle - \eta(\mathbf{x})), \quad (24)$$

where η denotes the phase of $f * \overline{\psi}$. Lemma 1 introduced low-frequency functions f_0 , with slow variations. Roughly speaking, since $\text{supp } f_0 \subset \mathcal{B}_{\infty}(\varepsilon/2)$, we can define a ‘‘minimal wavelength’’ $\lambda_{f_0} := 2\pi/\varepsilon$. Then,

$$\|\mathbf{h}\|_2 \ll \frac{2\pi}{\varepsilon} \implies f_0(\mathbf{x} + \mathbf{h}) \approx f_0(\mathbf{x}), \quad (25)$$

which leads to

$$(f * \text{Re}\overline{\psi})(\mathbf{x} + \mathbf{h}) \approx |(f * \overline{\psi})(\mathbf{x})| g(\mathbf{x}, \mathbf{h}). \quad (26)$$

On the one hand, we consider a continuous equivalent of the $\mathbb{C}\text{GMod}$ operator $\mathcal{C}_m[W]$ as introduced in (19). Such an operator, denoted by $\mathcal{C}[\psi]$, is defined, for any $f \in L^2_{\mathbb{R}}(\mathbb{R}^2)$, by

$$\mathcal{C}[\psi](f) : \mathbf{x} \mapsto |(f * \overline{\psi})(\mathbf{x})|. \quad (27)$$

On the other hand, we consider the continuous counterpart of $\mathbb{R}\text{GPool}$ as introduced in (22). It is defined as the maximum value of $f * \text{Re}\overline{\psi}$ over a sliding spatial window of size $r > 0$. This is possible because f and $\text{Re}\overline{\psi}$ both belong to $L^2_{\mathbb{R}}(\mathbb{R}^2)$, and therefore $f * \text{Re}\overline{\psi}$ is continuous. Such an operator, denoted by $\mathcal{R}_r[\psi]$, is defined, for any $f \in L^2_{\mathbb{R}}(\mathbb{R}^2)$, by

$$\mathcal{R}_r[\psi](f) : \mathbf{x} \mapsto \max_{\|\mathbf{h}\|_{\infty} \leq r} (f * \text{Re}\overline{\psi})(\mathbf{x} + \mathbf{h}). \quad (28)$$

For the sake of concision, the parameter between square brackets is ignored from now on.

If $r \ll 2\pi/\varepsilon$, then (26) is valid for any $\mathbf{h} \in \mathcal{B}_{\infty}(r)$. Then, using (27) and (28), we get

$$r \ll 2\pi/\varepsilon \implies \mathcal{R}_r f(\mathbf{x}) \approx \mathcal{C}f(\mathbf{x}) \max_{\|\mathbf{h}\|_{\infty} \leq r} g(\mathbf{x}, \mathbf{h}). \quad (29)$$

Using the periodicity of g , we can show that, if $r \geq \frac{\pi}{\|\nu\|_2}$, then $\mathbf{h} \mapsto g(\mathbf{x}, \mathbf{h})$ reaches its maximum value (= 1) on $\mathcal{B}_{\infty}(r)$. We therefore get

$$\frac{\pi}{\|\nu\|_2} \leq r \ll \frac{2\pi}{\varepsilon} \implies \mathcal{R}_r f(\mathbf{x}) \approx \mathcal{C}f(\mathbf{x}). \quad (30)$$

An exact quantification of the above approximation remains an open question. In the current paper, it will be provided as a conjecture, for the discrete framework.

2) *Adaptation to Discrete 2D Sequences:* We consider an input image $X \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$, a subsampling factor $m \in \mathbb{N}^*$ and a grid half-size $q \in \mathbb{N}^*$. We seek a relationship between

$$Y_{\text{pool}} := \mathcal{R}_{m,q}X \quad \text{and} \quad Y_{\text{mod}} := \mathcal{C}_{2m}X, \quad (31)$$

where \mathcal{C}_{2m} and $\mathcal{R}_{m,q}$ have been defined in (19) and (22), respectively. Note that, since max pooling also performs subsampling, both $\mathbb{C}\text{GMod}$ and $\mathbb{R}\text{GPool}$ operators, as defined in (31), have a subsampling factor equal to $2m$.

We now use the sampling results obtained in section III-C. Let f_X and $\psi_W \in \mathcal{U}_s$ denote the functions satisfying (11). On the one hand, we apply (13) in Proposition 2 to Y_{mod} . For any $\mathbf{n} \in \mathbb{Z}^2$,

$$\mathcal{C}_{2m}X[\mathbf{n}] = \mathcal{C}f_X(\mathbf{x}_n), \quad (32)$$

where $\mathbf{x}_n := 2ms\mathbf{n}$. On the other hand, we postulate that

$$\mathcal{R}_{m,q}X[\mathbf{n}] = \mathcal{R}_r f_X(\mathbf{x}_n) \quad (33)$$

for a certain value of $r \in \mathbb{R}_+^*$. Then, (30) implies $Y_{\text{mod}} \approx Y_{\text{pool}}$. However, as shown below, (33) is not satisfied. According to (22) and (23), we have

$$\mathcal{R}_{m,q}X[\mathbf{n}] = \max_{\|\mathbf{k}\|_\infty \leq q} \text{Re} \left((X * \overline{W}) \downarrow m \right)[2\mathbf{n} + \mathbf{k}]. \quad (34)$$

Therefore, according to (13) in Proposition 2, we get

$$\mathcal{R}_{m,q}X[\mathbf{n}] = \max_{\|\mathbf{k}\|_\infty \leq q} (f_X * \text{Re} \overline{\psi}_W)(\mathbf{x}_n + \mathbf{h}_k), \quad (35)$$

with

$$\mathbf{x}_n := 2ms\mathbf{n} \quad \text{and} \quad \mathbf{h}_k := msk. \quad (36)$$

By considering $r_q := ms \left(q + \frac{1}{2} \right)$, we get a variant of (33) in which the maximum is evaluated on a discrete grid of $(2q + 1)^2$ elements, instead of the continuous region $\mathcal{B}_\infty(r_q)$. As a consequence, (29) is replaced in the discrete framework by

$$q \ll 2\pi/(m\kappa) \implies \mathcal{R}_{m,q}X[\mathbf{n}] \approx \mathcal{C}_{2m}X[\mathbf{n}] \max_{\|\mathbf{k}\|_\infty \leq q} g_X(\mathbf{x}_n, \mathbf{h}_k), \quad (37)$$

with

$$g_X : (\mathbf{x}, \mathbf{h}) \mapsto \cos((\boldsymbol{\nu}, \mathbf{h}) - \eta_X(\mathbf{x})), \quad (38)$$

with $\boldsymbol{\nu} := \boldsymbol{\xi}/s$, and where η_X denotes the phase of $f_X * \overline{\psi}_W$. Unlike the continuous case, even if the window size r_q is large enough, the existence of $\mathbf{k} \in \{-q \dots q\}^2$ such that $g_X(\mathbf{x}_n, \mathbf{h}_k) = 1$ is not guaranteed, as illustrated in Fig. 2 with $q = 1$. Instead, we can only seek a probabilistic estimation of the relative quadratic error between Y_{pool} and Y_{mod} .

Approximation (37) implies

$$q \ll 2\pi/(m\kappa) \implies \|\mathcal{C}_{2m}X - \mathcal{R}_{m,q}X\|_2 \approx \|\delta_{m,q}X\|_2, \quad (39)$$

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

$$\delta_{m,q}X[\mathbf{n}] := \mathcal{C}_{2m}X[\mathbf{n}] \left(1 - \max_{\|\mathbf{k}\|_\infty \leq q} g_X(\mathbf{x}_n, \mathbf{h}_k) \right). \quad (40)$$

Expression (39) suggests that the ratio between the left and right terms can be bounded by a quantity which only depends on the product $m\kappa$ (subsampling factor \times frequency localization) and the grid half-size q :

$$\|\mathcal{C}_{2m}X - \mathcal{R}_{m,q}X\|_2 \leq (1 + \beta_q(m\kappa)) \|\delta_{m,q}X\|_2, \quad (41)$$

for some function $\beta_q : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ to be characterized. This is the goal of the following conjecture.

Conjecture 1: There exists

$$\beta_q : \mathbb{R}_+ \rightarrow \mathbb{R}_+ \quad \text{satisfying} \quad \beta_q(t) = O(t), \quad (42)$$

independent from the characteristic frequency $\boldsymbol{\xi} \in [-\pi, \pi]^2$, such that, for any $X \in \mathcal{L}_\mathbb{R}^2(\mathbb{Z}^2)$, (41) is satisfied.

We now seek a probabilistic bound for $\|\delta_{m,q}X\|_2$.




Fig. 2. Search for the maximum value of $\mathbf{h} \mapsto g_X(\mathbf{x}, \mathbf{h})$ over a discrete grid of size 3×3 , i.e., $q = 1$. This figure displays 3 examples with different frequencies $\boldsymbol{\nu} := \boldsymbol{\xi}/s$ and phases $\eta_X(\mathbf{x})$. Hopefully the result will be close to the true maximum (left), but there are some pathological cases in which all points in the grid fall into pits (middle and right).

E. Probabilistic Framework

1) *Notations:* In what follows, for any $z \in \mathbb{C}^*$, we denote by $\angle z \in [0, 2\pi[$ the argument of z . We now consider the unit circle $\mathbb{S}^1 \subset \mathbb{C}$. 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 going from z to z' counterclockwise:

$$[z, z']_{\mathbb{S}^1} := \{z'' \in \mathbb{S}^1 \mid \angle(z^*z'') \leq \angle(z^*z')\}. \quad (43)$$

By using the relation $\cos \alpha = \text{Re}(e^{i\alpha})$, (36) and (38) yield, for any $\mathbf{n} \in \mathbb{Z}^2$ and any $\mathbf{k} \in \{-q \dots q\}^2$,

$$g_X(\mathbf{x}_n, \mathbf{h}_k) = \text{Re}(z_X(\mathbf{x}_n)^* z_k), \quad (44)$$

where we have defined

$$z_X(\mathbf{x}) := e^{i\eta_X(\mathbf{x})} \quad \text{and} \quad z_k := e^{i(\boldsymbol{\nu}, \mathbf{h}_k)} = e^{im(\boldsymbol{\xi}, \mathbf{k})}. \quad (45)$$

Let us denote $N_q := (2q + 1)^2$. We consider a sequence, denoted by $(z_i^{(q)})_{i \in \{0 \dots N_q - 1\}}$, obtained by ordering $\{z_k\}_{k \in \{-q \dots q\}^2}$ in ascending order of their arguments:

$$0 = \theta_0^{(q)} \leq \dots \leq \theta_{N_q - 1}^{(q)} < 2\pi, \quad (46)$$

where we have denoted $\theta_i^{(q)} := \angle(z_i^{(q)})$. Besides, we extend the notations with $\theta_{N_q}^{(q)} := 2\pi$ and $z_{N_q}^{(q)} := z_0^{(q)}$. Then, we split \mathbb{S}^1 into N_q arcs delimited by $(z_i^{(q)})_{i \in \{0 \dots N_q - 1\}}$:

$$\mathfrak{A}_i^{(q)} := \begin{cases} [z_i^{(q)}, z_{i+1}^{(q)}]_{\mathbb{S}^1} & \text{if } \theta_{i+1}^{(q)} - \theta_i^{(q)} < 2\pi; \\ \mathbb{S}^1 & \text{otherwise.} \end{cases} \quad (47)$$

Finally, for any $i \in \{0 \dots N_q - 1\}$, we denote by $\omega_i^{(q)} := \theta_{i+1}^{(q)} - \theta_i^{(q)}$ the angular measure of arc $\mathfrak{A}_i^{(q)}$.

Remark 2: According to (45), the above quantities depend on the product $m \times \boldsymbol{\xi} \in \mathbb{R}^2$. Therefore, we will sometimes write $\omega_i^{(q)}(m\boldsymbol{\xi})$, where $\omega_i^{(q)}$ is defined as a function of \mathbb{R}^2 :

$$\omega_i^{(q)} : \mathbb{R}^2 \rightarrow [0, 2\pi]. \quad (48)$$

2) *Random Variables*: From now on, input X is considered as discrete 2D stochastic processes. In order to “randomize” f_X introduced in (11), we define a continuous stochastic process from X , denoted by F_X , such that

$$\forall \mathbf{x} \in \mathbb{R}^2, F_X(\mathbf{x}) := \sum_{\mathbf{n} \in \mathbb{Z}^2} X[\mathbf{n}] \phi_{\mathbf{n}}^{(s)}(\mathbf{x}). \quad (49)$$

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

$$M_X := |F_X * \bar{\psi}|; \quad H_X := \angle(F_X * \bar{\psi}); \quad Z_X := e^{iH_X}, \quad (50)$$

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

$$G_{X, \mathbf{k}} := \text{Re}(Z_X^* z_{\mathbf{k}}); \quad G_{\max X} := \max_{\|\mathbf{k}\|_{\infty} \leq q} G_{X, \mathbf{k}}. \quad (51)$$

For any $\mathbf{x} \in \mathbb{R}^2$, $f_X(\mathbf{x})$ and $\eta_X(\mathbf{x})$, such as introduced in (11) and (38), are respectively drawn from $F_X(\mathbf{x})$ and $H_X(\mathbf{x})$. Then, $z_X(\mathbf{x})$ such as introduced in (45) is a realization of $Z_X(\mathbf{x})$. Consequently, according to (44), $g_X(\mathbf{x}, \mathbf{h}_{\mathbf{k}})$ is a realization of $G_{X, \mathbf{k}}(\mathbf{x})$. Besides, according to the definition of $\mathbb{C}\text{GMod}$ in (19) and \mathbf{x}_n in (36), Proposition 2 implies that

$$M_X(\mathbf{x}_n) = \mathcal{C}_{2m} X[\mathbf{n}]. \quad (52)$$

We remind that $\boldsymbol{\xi} \in [-\pi, \pi]^2$ and $\kappa \in]0, 2\pi]$ respectively denote the center and size of the Fourier support of W , as introduced in section III-C. To compute the expected discrepancy between Y_{pool} and Y_{mod} , we assume that

$$\|\boldsymbol{\xi}\|_2 \gg 2\pi/M; \quad (53)$$

$$\|\boldsymbol{\xi}\|_2 \gg \kappa, \quad (54)$$

where $M \in \mathbb{N}^*$ 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 H.

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

Hypothesis 2: For any $n \in \mathbb{N}^*$ and $\mathbf{x}, \mathbf{y}_1, \dots, \mathbf{y}_n \in \mathbb{R}^2$, the random variables $M_X(\mathbf{y}_i)$ for $i \in \{1 \dots n\}$ are jointly independent of $Z_X(\mathbf{x})$.

F. Expected Quadratic Error between $\mathbb{R}\text{GPool}$ and $\mathbb{C}\text{GMod}$

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

$$\tilde{P}_X := \|\mathcal{C}_{2m} X - \mathcal{R}_{m, q} X\|_2 / \|\mathcal{C}_{2m} X\|_2. \quad (55)$$

According to (31), this is an estimation of the relative error between Y_{mod} and Y_{pool} .

First, let us reformulate $\delta_{m, q} X$, introduced in (40), using the probabilistic framework. According to (44) and (51), we have, for any $\mathbf{n} \in \mathbb{Z}^2$,

$$\delta_{m, q} X[\mathbf{n}] := \mathcal{C}_{2m} X[\mathbf{n}] (1 - G_{\max X}(\mathbf{x}_n)). \quad (56)$$

We now consider the stochastic process $Q_X := 1 - G_{\max X}$, and the random variable

$$\tilde{Q}_X := \|\delta_{m, q} X\|_2 / \|\mathcal{C}_{2m} X\|_2. \quad (57)$$

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

The first point is established in Proposition 4 below, and the two remaining ones are the purpose of Theorem 2.

Proposition 4: Assuming Hypothesis 1, the expected value of $Q_X(\mathbf{x})^2$ is independent from the choice of $\mathbf{x} \in \mathbb{R}^2$, and

$$\mathbb{E}[Q_X(\mathbf{x})^2] = \gamma_q(m\boldsymbol{\xi})^2, \quad (58)$$

where we have defined

$$\gamma_q : \zeta \mapsto \sqrt{\frac{3}{2} + \frac{1}{4\pi} \sum_{i=0}^{N_q-1} \left(\sin \omega_i^{(q)}(\zeta) - 8 \sin \frac{\omega_i^{(q)}(\zeta)}{2} \right)}, \quad (59)$$

with $\omega_i^{(q)} : \mathbb{R}^2 \mapsto [0, 2\pi]$ being introduced in (48).

Proof: See Appendix J. ■

We consider an ideal scenario where the $z_{\mathbf{k}}$ are evenly spaced on \mathbb{S}^1 . Then, an order 2 Taylor expansion yields $\gamma_q(\zeta) = o(1/q^2)$, meaning that \tilde{Q}_X quickly vanishes when the grid half-size q increases. Fig. 4 displays $\boldsymbol{\xi} \mapsto \gamma_q(m\boldsymbol{\xi})^2$ for $\boldsymbol{\xi} \in [-\pi, \pi]^2$, with $m = 4$ and $q = 1$ as in AlexNet. We notice that, for the major part of the Fourier domain, γ_q remains close to 0. However, we observe a regular pattern of dark regions, which correspond to pathological frequencies where the repartition of $z_{\mathbf{k}}$ is unbalanced.

So far, we established a result at the pixel level. Before stating Theorem 2, we need the following intermediate result.

Proposition 5: We consider the random variable

$$\tilde{S}_X := \|\mathcal{C}_{2m} X\|_2. \quad (60)$$

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

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

Proof: See Appendix K. ■

Finally, Propositions 4 and 5 yield the following theorem. It provides an upper bound on the expected value of the relative quadratic error \tilde{P}_X^2 , such as defined in (55).

Theorem 2: We assume that Conjecture 1 is true. Then, under Hypotheses 1 and 2, we have

$$\mathbb{E}[\tilde{P}_X^2] \leq (1 + \beta_q(m\kappa))^2 \gamma_q(m\boldsymbol{\xi})^2, \quad (61)$$

with β_q and γ_q being introduced in (42) and (59), respectively.

Proof: See Appendix L. ■

G. Shift Invariance of $\mathbb{R}\text{GPool}$ Outputs

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

Theorem 3: Let $W \in \mathcal{G}(\boldsymbol{\xi}, \kappa)$ denote a discrete Gabor-like filter, $m \in \mathbb{N}^*$ a subsampling factor and $q \in \mathbb{N}^*$ a grid half-size. We consider a stochastic process X whose realizations

are elements of $l_{\mathbb{R}}^2(\mathbb{Z}^2)$. We assume that Hypotheses 1 & 2 are satisfied.⁸ Besides, we assume Conjecture 1.

Given a translation vector $\mathbf{u} \in \mathbb{R}^2$, we consider the following random variable:

$$\tilde{R}_{X, \mathbf{u}} := \|\mathcal{R}_{m, q}(\mathcal{T}_{\mathbf{u}}X) - \mathcal{R}_{m, q}X\|_2 / \|\mathcal{C}_{2m}X\|_2. \quad (62)$$

Then, under the following conditions:

$$\kappa \leq \pi/m \quad \text{and} \quad \|\mathbf{u}\|_1 \leq \pi/\kappa, \quad (63)$$

we have

$$\mathbb{E}[\tilde{R}_{X, \mathbf{u}}] \leq 2(1 + \beta_q(m\kappa))\gamma_q(m\xi) + \alpha(\kappa\mathbf{u}), \quad (64)$$

where α , β_q and γ_q are defined in (8), (42) and (59).

Proof: See Appendix M. ■

If κ is sufficiently small, then $\alpha(\kappa\mathbf{u})$ and $\beta_q(m\kappa)$ become negligible with respect to $\gamma_q(m\xi)$, and the bound given in (64) is roughly equal to $2\gamma_q(m\xi)$. Theorem 3 therefore provides a validity domain for shift invariance of $\mathbb{R}\text{GPool}$ operators, as illustrated in Fig. 4 with $q = 1$.

Remark 3: The stochastic discrepancy introduced in (62) is estimated relatively to the $\mathbb{C}\text{GMod}$ output. This choice is motivated by the perfect shift invariance of its norm, as shown in Proposition 3.

Remark 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 concision, in the remaining of this paper, we will drop q in the notations, which implicitly means $q = 1$.

IV. STANDARD CNNs IN OUR FRAMEWORK

In this section, we show how the theoretical results established above on single-channel inputs can be applied to standard CNNs such as AlexNet or ResNet, which are usually designed for multichannel inputs—e.g., RGB images.

A. Background

We focus on the first layers of a classical CNN architecture, in which input images are propagated through a convolution layer followed by rectified linear unit and max pooling operators, as described below.

We denote by K and $L \in \mathbb{N}^*$ the number of input and output channels, respectively. The convolution layer is characterized by a trainable weight tensor $\mathbf{V} \in (l_{\mathbb{R}}^2(\mathbb{Z}^2))^{L \times K}$ with a finite support, a bias vector $\mathbf{b} \in \mathbb{R}^L$ and a subsampling factor $m \in \mathbb{N}^*$. Considering an input image $\mathbf{X} \in (l_{\mathbb{R}}^2(\mathbb{Z}^2))^K$ with finite support, we denote by $\mathbf{A} \in (l_{\mathbb{R}}^2(\mathbb{Z}^2))^L$ the output of the max pooling layer. Then we have, for any $l \in \{0 \dots L - 1\}$,

$$A_l := (\text{MaxPool} \circ \text{ReLU}) \left(b_l + \sum_{k=0}^{K-1} (X_k * \bar{V}_{lk}) \downarrow m \right), \quad (65)$$

where MaxPool has been introduced in (23) with $q = 1$ (see Remark 4), and ReLU is defined such that, for any $X' \in l_{\mathbb{R}}^2(\mathbb{Z}^2)$ and any $\mathbf{n} \in \mathbb{Z}^2$, $\text{ReLU}(X')[\mathbf{n}] := \max(0, X'[\mathbf{n}])$.

⁸We can easily prove that these properties are independent from the choice of sampling interval $s > 0$.

Expression (65) also introduces a bias notation, defined such that $(b + X')[\mathbf{n}] = b + X'[\mathbf{n}]$ for any $\mathbf{n} \in \mathbb{Z}^2$.

In many cases, input images are composed of three RGB channels; therefore $K = 3$. The other parameters depend on the chosen CNN architecture. For example, in AlexNet, the weight tensor \mathbf{V} is supported in a region of size 11×11 and the subsampling factor m is equal to 4. ResNet models use kernels of size 7×7 with $m = 2$.

B. Making CNNs Compatible with our Theoretical Results

The bias and ReLU are outside the scope of our study. However, we can show that $A_l = \text{ReLU}(b_l + Y_{\text{pool}l})$, where we have defined $Y_{\text{pool}l} := \mathcal{R}_m^{(l)}\mathbf{X}$, with

$$\mathcal{R}_m^{(l)} : \mathbf{X} \mapsto \text{MaxPool} \left(\sum_{k=0}^{K-1} (X_k * \bar{V}_{lk}) \downarrow m \right). \quad (66)$$

For the sake of our study, we therefore consider a strictly equivalent CNN architecture where the bias and ReLU are computed after the max pooling layer. We then focus on the intermediate output $\mathbf{Y}_{\text{pool}} \in (l_{\mathbb{R}}^2(\mathbb{Z}^2))^L$.

Remark 5: In many architectures including ResNet, bias \mathbf{b} is replaced by a batch normalization layer with affine transformation [49]. Swapping such a layer with max pooling isn't straightforward, but can nevertheless be done with caution. Therefore, we can once again focus on \mathbf{Y}_{pool} such as introduced in (66).

In what follows, we assume that the network has been trained on a natural image dataset such as ImageNet. Let \mathcal{L} denote the subset of output channels $l \in \{0 \dots L - 1\}$ such that, for any $k \in \{0 \dots K - 1\}$, V_{lk} behaves like a band-pass filter (see Fig. 1). They are referred to as *Gabor channels*. We would like to apply the theoretical results obtained in section III on $Y_{\text{pool}l}$, for any $l \in \mathcal{L}$.

To do so, we need to show that $Y_{\text{pool}l}$ is the output of a $\mathbb{R}\text{GPool}$ operator such as introduced in (22), for some weight and input to be defined. This is in general not possible though, because the sum operator in (66) cannot be interchanged with max pooling. To solve this problem, we state the following hypothesis, applied on any Gabor channel $l \in \mathcal{L}$. It states that the trained kernels V_{lk} for $k \in \{0 \dots K - 1\}$ are identical, up to a multiplicative constant.

Hypothesis 3: Let $\tilde{V}_l := \frac{1}{K} \sum_{k=0}^{K-1} V_{lk}$ denote the mean kernel of the l -th output channel. Then, for any $l \in \mathcal{L}$, there exists $\mu_l \in \mathbb{R}^K$ such that

$$\forall k \in \{0 \dots K - 1\}, V_{lk} = \mu_{lk} \tilde{V}_l. \quad (67)$$

Intuitively, when looking at Gabor-like kernels in Fig. 1, they roughly appear grayscale. This observation supports Hypothesis 3 with $\mu_{lk} \approx 1$. A more accurate justification for this hypothesis is provided in Appendix I.

Then, considering the linear combination $\tilde{X}_l := \mu_l^\top \mathbf{X}$, we apply Hypothesis 3 to (66):

$$\forall l \in \mathcal{L}, Y_{\text{pool}l} = \mathcal{R}_m^{(l)}\mathbf{X} = \mathcal{R}_m[\tilde{W}_l](\tilde{X}_l), \quad (68)$$

TABLE I
PERCENTAGE OF ENERGY WITHIN A FOURIER WINDOW OF SIZE $\kappa \times \kappa$,
FOR GABOR-LIKE FILTERS $\{\widetilde{W}_l\}_{l \in \mathcal{L}}$

MODEL	NB CHANNELS	SIZE OF \mathcal{L}	κ	MEAN RATIO
ALEXNET	64	26	$\pi/4$	67%
RESNET-34	64	22	$\pi/2$	76%

where $\widetilde{W}_l \in l^2(\mathbb{Z}^2)$ denotes the complex-valued companion of \widetilde{V}_l satisfying (5). We also define $\mathcal{C}_{2m}^{(l)}$ and its associated output feature map $Y_{\text{mod}l}$, such that

$$\forall l \in \mathcal{L}, Y_{\text{mod}l} := \mathcal{C}_{2m}^{(l)} \mathbf{X} = \mathcal{C}_{2m}[\widetilde{W}_l](\widetilde{X}_l). \quad (69)$$

Besides, the following hypothesis states that, for any $l \in \mathcal{L}$, \widetilde{W}_l is a discrete Gabor-like filter for which the Fourier support size is shared among the Gabor channels.

Hypothesis 4: There exists $\kappa \leq \pi/m$ such that, for any Gabor channel $l \in \mathcal{L}$,

$$\exists \xi_l \in [-\pi, \pi]^2 : \widetilde{W}_l \in \mathcal{G}(\xi_l, \kappa). \quad (70)$$

Let $l \in \mathcal{L}$. According to Hypotheses 3–4, $\mathcal{R}_m^{(l)}$ and $\mathcal{C}_{2m}^{(l)}$ such as defined in (68) and (69) can be respectively qualified as $\mathbb{R}\text{GPool}$ and $\mathbb{C}\text{GMod}$ operators. We now consider

$$\widetilde{P}_{\mathbf{X}}^{(l)} := \|\mathcal{C}_{2m}^{(l)} \mathbf{X} - \mathcal{R}_m^{(l)} \mathbf{X}\|_2 / \|\mathcal{C}_{2m}^{(l)} \mathbf{X}\|_2; \quad (71)$$

$$\widetilde{R}_{\mathbf{X}, \mathbf{u}}^{(l)} := \|\mathcal{R}_m^{(l)}(\mathcal{T}_{\mathbf{u}} \mathbf{X}) - \mathcal{R}_m^{(l)} \mathbf{X}\|_2 / \|\mathcal{C}_{2m}^{(l)} \mathbf{X}\|_2. \quad (72)$$

Then, under Hypotheses 1–4, we can apply Theorems 2 and 3 to the above random variables, thus providing a shift invariance measure for a subset of outputs in CNNs:

$$\mathbb{E} \left[\widetilde{P}_{\mathbf{X}}^{(l)^2} \right] \leq (1 + \beta(m\kappa))^2 \gamma(m\xi_l)^2; \quad (73)$$

$$\mathbb{E} \left[\widetilde{R}_{\mathbf{X}, \mathbf{u}}^{(l)} \right] \leq 2(1 + \beta(m\kappa)) \gamma(m\xi_l) + \alpha(\kappa\mathbf{u}), \quad (74)$$

where ξ_l has been introduced in (70).

In practice, Hypothesis 4 cannot be exactly satisfied. This is because \widetilde{V}_l is finitely supported, and thus its power spectrum cannot be exactly zero on a region with non-zero measure. To evaluate how close we are to this ideal situation, we measured the maximum percentage of energy within a window of size $\kappa \times \kappa$ in the Fourier domain, with respect to the whole filter W_l . We then computed the mean percentage over all the Gabor channels $l \in \mathcal{L}$. The results are shown in Table I, for AlexNet and ResNet after training with ImageNet. The window size κ has been set to its highest admissible value, i.e., π/m .

We notice that residual energy outside the window of interest is quite high, especially for AlexNet. Therefore, W_l deviates from a “perfect” Gabor-like filter, which may lead to higher instabilities.

V. OPERATORS BASED ON THE DUAL-TREE WAVELET PACKET TRANSFORM

In order to validate the results established in section III, we consider the dual-tree wavelet packet transform (DT-CWPT), a fast algorithm which achieves subsampled convolutions with discrete Gabor-like filters. As explained below, DT-CWPT

spawns a set of filters which tiles the whole frequency domain in a regular fashion. Furthermore, increasing the subsampling factor m results in a decreased Fourier support size $\kappa = \pi/m$, therefore matching the condition stated in (63). This is also consistent with what is observed in CNNs such as AlexNet or ResNet. DT-CWPT thus provides a convenient framework to emulate the behavior of an actual CNN while testing Theorems 1–3 in a controlled environment.

A. Discrete Wavelet Packet Transform

This is a brief overview on the classical, real-valued 2D wavelet packet transform (WPT) algorithm [48, p. 377], which is the starting point for building the redundant, complex-valued DT-CWPT.

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

$$G_0 = h \otimes h; G_1 = h \otimes g; G_2 = g \otimes h; G_3 = g \otimes g. \quad (75)$$

Let $X \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$. The decomposition starts with $X_0^{(0)} = X$. Given $j \in \mathbb{N}$, suppose that we have computed 4^j sequences of wavelet packet coefficients at stage j , denoted by $X_l^{(j)} \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$ 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 $X_l^{(j)}$ into four sub-sequences, using subsampled (or strided) convolutions with kernels G_k , for each $k \in \{0..3\}$:

$$\forall k \in \{0..3\}, X_{4l+k}^{(j+1)} = \left(X_l^{(j)} * \overline{G_k} \right) \downarrow 2. \quad (76)$$

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

$$\mathbf{X}^{(J)} := \left(X_l^{(J)} \right)_{l \in \{0..4^J - 1\}} \quad (77)$$

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

The following proposition introduces an array of resulting kernels $\mathbf{V}^{(J)}$ which is illustrated in Fig. 3 with $J = 2$.

Proposition 6: For any $l \in \{0..4^J - 1\}$, there exists $V_l^{(J)} \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$ such that

$$X_l^{(J)} = \left(X * \overline{V_l^{(J)}} \right) \downarrow 2^J. \quad (78)$$

Proof: See Appendix N. ■

B. 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, N. Kingsbury designed a new type of discrete wavelet transform [29], where images are decomposed in a redundant frame of nearly-analytic, complex-valued waveforms. It was




Fig. 3. Resulting kernels $\mathbf{V}^{(2)}$ for WPT (left) and $\mathbf{W}^{(2)}$ for DT-CWPT (right), computed with Q-shift orthogonal QMFs of length 10 [50]. They have been cropped to size 11×11 for the sake of visibility. The right part of the figure displays 32 complex filters, alternatively represented by their real and imaginary parts. The feature maps related to ① and ② are obtained with two distinct formulas, which are summarized in (82). Illustration from [31].

later extended to the wavelet packet framework [30]. The latter operation, referred to as *dual-tree complex wavelet packet transform* (DT-CWPT), is performed as follows.

Let (h_0, g_0) and (h_1, g_1) denote two pairs of QMFs as defined in section V-A, satisfying the *half-sample delay* condition:

$$\forall \omega \in [-\pi, \pi], \hat{h}_1(\omega) = e^{-i\omega/2} \hat{h}_0(\omega). \quad (79)$$

Then, for any $k \in \{0..3\}$, we build a two-dimensional FB $\mathbf{G}_k := (\mathbf{G}_{k,l})_{l \in \{0..3\}}$ similarly to (75):

$$\mathbf{G}_{k,0} = \mathbf{h}_i \otimes \mathbf{h}_j; \quad \mathbf{G}_{k,1} = \mathbf{h}_i \otimes \mathbf{g}_j; \quad (80)$$

$$\mathbf{G}_{k,2} = \mathbf{g}_i \otimes \mathbf{h}_j; \quad \mathbf{G}_{k,3} = \mathbf{g}_i \otimes \mathbf{g}_j, \quad (81)$$

where $i, j \in \{0, 1\}$ are defined such that $k = 2 \times i + j$.⁹

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 \mathbf{X} into four multichannel WPT representations $\mathbf{X}_{0-3}^{(J)}$, each of which satisfies (76) and (77). Then, for any $l \in \{0..4^J - 1\}$, the following complex feature maps are computed:

$$\begin{pmatrix} Z_l^{(J)} \\ Z_{4^J+l}^{(J)} \end{pmatrix} = \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} X_{0,l}^{(J)} \\ X_{3,l}^{(J)} \end{pmatrix} + i \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} X_{2,l}^{(J)} \\ X_{1,l}^{(J)} \end{pmatrix}. \quad (82)$$

We denote by $L_J := 2 \cdot 4^J$ the number of output feature maps. Then,

$$\mathbf{Z}^{(J)} := (Z_l^{(J)})_{l \in \{0..L_J-1\}} \quad (83)$$

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

As for standard WPT, the following proposition introduces an array of resulting (complex-valued) kernels $\mathbf{W}^{(J)}$ for which a graphical representation is provided in Fig. 3 with $J = 2$.

⁹Actually, the FB design requires some technicalities which are not described here.

Proposition 7: For any $l \in \{0..L_J - 1\}$, there exists $W_l^{(J)} \in l^2(\mathbb{Z}^2)$ such that

$$Z_l^{(J)} = \left(\mathbf{X} * \overline{W_l^{(J)}} \right) \downarrow 2^J. \quad (84)$$

Proof: See Appendix O. ■

C. Invariance Results Applied to the Dual-Tree Framework

We assume that h_0 is a Shannon filter, such that $\hat{h}_0(\omega) := \sqrt{2}$ if $\omega \in [-\frac{\pi}{2}, \frac{\pi}{2}]$ and 0 otherwise. Let $J \in \mathbb{N}^*$ denote the number of decomposition stages. We consider the resulting kernels $W_l^{(J)}$ satisfying (84). The following hypothesis states that DT-CWPT tiles the frequency plane with a square grid.

Hypothesis 5: For any $l \in \{0..L_J - 1\}$, there exists $\sigma_l^{(J)} \in \{-2^J..2^J - 1\}^2$ such that

$$W_l^{(J)} \in \mathcal{G}(\xi_l^{(J)}, \kappa_J), \quad (85)$$

where we have defined

$$\xi_l^{(J)} := \left(\sigma_l^{(J)} + \frac{1}{2} \right) \frac{\pi}{2^J} \quad \text{and} \quad \kappa_J := \frac{\pi}{2^J}. \quad (86)$$

It can be shown that Hypothesis 5 is an asymptotic result, when J goes to ∞ . In reality, the Fourier support of $W_l^{(J)}$ is contained in four symmetric square regions of size κ_J . If the dual filter h_1 satisfies the half-sample delay condition (79), then the energy of $W_l^{(J)}$ goes to 0 in all but one of the four regions (relatively to the filter's total energy). We nevertheless consider Hypothesis 5 as reasonable when $J \geq 2$.¹⁰

Then, according to (84) and (85), we can apply Theorems 1–3 to the dual-tree framework. More precisely, for any output channel $l \in \{0..L_J - 1\}$, we consider

$$Y_{\text{pool}}^{(J)} := \mathcal{R}_l^{(J)} \mathbf{X} \quad \text{and} \quad Y_{\text{mod}}^{(J)} := \mathcal{C}_l^{(J)} \mathbf{X}, \quad (87)$$

where we have denoted, for the sake of readability,

$$\mathcal{R}_l^{(J)} := \mathcal{R}_{m_J} [W_l^{(J)}] \quad \text{and} \quad \mathcal{C}_l^{(J)} := \mathcal{C}_{2m_J} [W_l^{(J)}], \quad (88)$$

with $m_J := 2^{J-1}$. We recall that \mathcal{R}_m and \mathcal{C}_{2m} have been introduced in (31) with $q \leftarrow 1$, for any $m \in \mathbb{N}^*$.

Both $\mathcal{R}_l^{(J)}$ and $\mathcal{C}_l^{(J)}$ perform DT-CWPT with J decomposition stages. However, the max pooling operator requires an extra level of subsampling. To counterbalance this, when computing $Y_{\text{pool}}^{(J)}$, the last stage of DT-CWPT decomposition must be performed without subsampling, at the cost of increased redundancy.¹¹ This is why $m_J := 2^{J-1}$ instead of 2^J .

In order Theorems 1 and 3 to be applicable, we need to check that conditions stated in (63) are satisfied. Using the DT-CWPT framework, they become

$$\kappa_J \leq \frac{\pi}{2^{J-1}} \quad \text{and} \quad \|\mathbf{u}\|_1 \leq \frac{\pi}{\kappa_J} = 2^J. \quad (89)$$

The first condition is always met, according to (86). As for the second one, it establishes a limit on $\|\mathbf{u}\|_1$ above which

¹⁰This asymptotic result is not true for “edge filters”, i.e., when $\|\xi_l^{(J)}\|_\infty = (1 - 2^{-(J+1)})\pi$. In this case, a small fraction of the filter's energy remains located at the far end of the Fourier domain [30]. However, this edge effect is ignored and we still consider Hypothesis 5 as reasonable.

¹¹This is similar to the concept of stationary wavelet transform [51].

shift invariance can no longer be estimated. Note however that shifting the input by 2^J pixels results in a 1-pixel output shift. Therefore, $\mathcal{R}_l^{(J)}(\mathcal{T}_u X)$ can always be compared with a shifted version of $\mathcal{R}_l^{(J)}X$. We then get a partial measure of shift equivariance.

In the Shannon setting, $h_0[n]$, $g_0[n]$, $h_1[n]$ and $g_1[n]$ decay in $O(1/n)$, which makes them difficult to approximate in practice. It requires very large vectors to avoid numerical instabilities. Practical implementations use fast-decaying filters such as Meyer QMFs [52], or finite-length filters which approximate the half-sample delay condition [53]. Therefore, residual energy can be observed outside the Fourier windows introduced in Hypothesis 5. To counterbalance this, we relax this hypothesis by increasing the window size up to $\kappa_{J} := \pi/2^{J-1}$, which is closer to what is observed in standard convolutions layers (see Table I).

VI. EXPERIMENTS AND RESULTS

In this section, we experimentally validate our theoretical results. To do so, we built operators based on DT-CWPT, as explained in section V. Using a dataset of natural images, we measured the mean discrepancy between $\mathbb{R}GPool$ and $\mathbb{C}GMod$ outputs, and evaluated the shift invariance of both models.

Dual-tree decompositions have been performed with Q-shift orthogonal filters of length 10 [50], which approximately meets the half-sample delay condition (79).

Our implementation is based on PyTorch. The models were evaluated on the validation set of ImageNet ILSVRC2012 [54], which contains $N_{\text{data}} := 50\,000$ images.

A. Discrepancies between $\mathbb{R}GPool$ and $\mathbb{C}GMod$

Each image $n \in \{0 \dots N_{\text{data}} - 1\}$ in the dataset was converted to grayscale, from which a center crop of size 224×224 was extracted. We denote by $X_n \in l_{\mathbb{R}}^2(\mathbb{Z}^2)$ the resulting input feature map. For any $l \in \{0 \dots L_J - 1\}$, we denote by $Y_{\text{pool}nl}^{(J)}$ and $Y_{\text{mod}nl}^{(J)}$ the outputs satisfying (87) with $X \leftarrow X_n$.

Then, the relative quadratic error between $Y_{\text{mod}nl}^{(J)}$ and $Y_{\text{pool}nl}^{(J)}$ was computed. It is defined by the square of

$$\rho_{nl}^{(J)} := \|Y_{\text{mod}nl}^{(J)} - Y_{\text{pool}nl}^{(J)}\|_2 / \|Y_{\text{mod}nl}^{(J)}\|_2. \quad (90)$$

Finally, for each output channel l , an empirical estimate for $\mathbb{E}[\tilde{P}_X^2]$ such as introduced in (55) was obtained by averaging $(\rho_{nl}^{(J)})^2$ over the whole dataset. We denote by $(\tilde{\rho}_l^{(J)})^2$ the corresponding quantity.

Since $\mathcal{R}_l^{(J)}$ and $\mathcal{C}_l^{(J)}$ are parameterized by $W_l^{(J)}$, the value of $(\tilde{\rho}_l^{(J)})^2$ depends on the filter's characteristic frequency $\xi_l^{(J)}$. According to Hypothesis 5, these frequencies form a regular grid in the Fourier domain. This provides a visual representation of $(\tilde{\rho}_l^{(J)})^2$, as shown in Fig. 5. We can observe a regular pattern of dark spots. More precisely, high discrepancies between max pooling and modulus seem to occur when the support of $W_l^{(J)}$ is located in a dark region of Fig. 4. This result corroborates the theoretical study, which states that high discrepancies are expected for certain pathological frequencies, due to the search for a maximum value over a discrete grid—see Fig. 2.




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




Fig. 5. Relative quadratic error between the outputs of $\mathbb{R}GPool$ and $\mathbb{C}GMod$. For each wavelet packet channel $l \in \{0 \dots L_J - 1\}$, $(\tilde{\rho}_l^{(J)})^2$ is represented as a grayscale pixel centered in $\xi_l^{(J)}$, such as introduced in (86). Since the subsampling factor m_J is equal to 2^{J-1} , these empirical estimates can be compared with the left and right parts of Fig. 4. The plots are symmetrized to account for the complex conjugate feature maps.

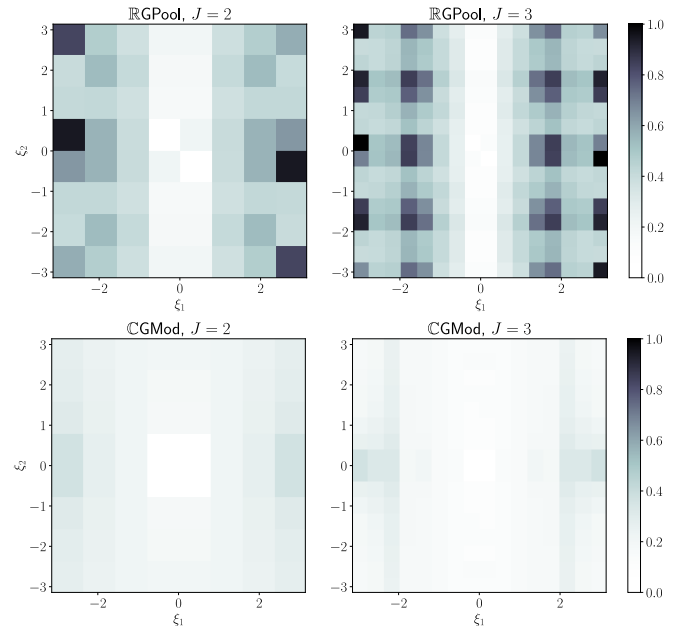


Fig. 6. Empirical measure of horizontal shift invariance of $\mathbb{R}GPool$ and $\mathbb{C}GMod$ outputs. For each $l \in \{0 \dots L_J - 1\}$, $\tilde{\rho}_{\text{pool}l}^{(J)}$ (top) and $\tilde{\rho}_{\text{mod}l}^{(J)}$ (bottom) are represented as a grayscale pixel centered in $\xi_l^{(J)}$, such as introduced in (86).

B. Shift invariance

For each input image X_n 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:

$$\rho_{\text{pool}nl}^{(J)} := \|Y_{\text{pool}nl}^{(J)} - Y_{\text{mod}nl}^{(J)}\|_2 / \|Y_{\text{mod}nl}^{(J)}\|_2, \quad (91)$$

where $Y_{\text{pool}}^{(J)}$, $Y_{\text{pool}}^{(J)}$ and $Y_{\text{mod}}^{(J)}$ satisfy (87) with $X \leftarrow X_n$ or X'_n . Finally, for each output channel $l \in \{0 \dots L_J - 1\}$, an empirical estimate for $\mathbb{E}[R_{X,u}]$ such as introduced in (62), with $u := (1, 0)$, was obtained by averaging $\rho_{\text{pool}nl}^{(J)}$ over the whole dataset. We denote by $\tilde{\rho}_{\text{pool}l}^{(J)}$ the corresponding quantity. We point out that shift invariance is measured relatively to the norm of the CGMod output, as explained in Remark 3.

On the other hand, the same procedure was applied on the CGMod operators:

$$\rho_{\text{mod}nl}^{(J)} := \|Y_{\text{mod}nl}^{(J)} - Y_{\text{mod}nl}^{(J)}\|_2 / \|Y_{\text{mod}nl}^{(J)}\|_2, \quad (92)$$

and $\tilde{\rho}_{\text{mod}l}^{(J)}$ was obtained as before by averaging $\rho_{\text{mod}nl}^{(J)}$ over the whole dataset.

A visual representation of $\tilde{\rho}_{\text{pool}l}^{(J)}$ and $\tilde{\rho}_{\text{mod}l}^{(J)}$ are provided in Fig. 6. Two observations can be drawn here.

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 (not shown in this paper).

Elsewhere, we observe that high discrepancies between $\mathbb{R}G\text{Pool}$ and $\mathbb{C}G\text{Mod}$ outputs (Fig. 5) are correlated with shift instability of $\mathbb{R}G\text{Pool}$ (Fig. 6, top). This is in line with (61) and (64) in Theorems 2–3. Note that $\mathbb{C}G\text{Mod}$ outputs are nearly shift invariant regardless the characteristic frequency $\xi_l^{(J)}$ (Fig. 6, bottom), as predicted by Theorem 1 (20).

VII. CONCLUSION

In this paper, we studied shift invariance properties captured by the max pooling operator, when applied on top of a convolution layer with Gabor-like kernels. More precisely, we established a validity domain for near-shift invariance, and confirmed our predictions with an experimental setting based on the dual-tree complex wavelet packet transform.

As shown in this paper, the $\mathbb{C}G\text{Mod}$ operator acts like a proxy for $\mathbb{R}G\text{Pool}$, extracting comparable features with higher stability. This result suggests a way to build an architecture sharing the structure and behavior of a standard network, except that shift invariance would be improved. This could be done by considering a DT-CWPT-based twin model such as introduced in our workshop paper [31], and replacing $\mathbb{R}G\text{Pool}$ by $\mathbb{C}G\text{Mod}$ operators as done above in a deterministic context.

Since CNNs generally implement successive blocks of convolution and pooling layers, another line of research would be to extend our results to a cascade of $\mathbb{R}G\text{Pool}$ operators. Our work is thus an important step towards deeper understanding of popular networks, based on the wavelet theory.

APPENDIX A PROOF OF LEMMA 1

Proof: We can show that, for any $\omega \in \mathbb{R}^2$,

$$\widehat{f}_0(\omega) = (\widehat{f * \psi})(\omega - \nu) = \mathcal{T}_\nu(\widehat{f\psi})(\omega).$$

By hypothesis on ψ , we have $\text{supp } \widehat{\psi} \subset \mathcal{B}_\infty(-\nu, \varepsilon/2)$, which yields the result. \blacksquare

APPENDIX B PROOF OF PROPOSITION 1

Proof: Using the 2D Plancherel formula, we compute

$$\begin{aligned} \|\mathcal{T}_h f_0 - f_0\|_{L^2}^2 &= \frac{1}{4\pi^2} \|\widehat{\mathcal{T}_h f_0} - \widehat{f_0}\|_{L^2}^2 \\ &= \frac{1}{4\pi^2} \iint_{\mathcal{B}_\infty(\varepsilon/2)} |\widehat{f_0}(\omega)|^2 |e^{-i\langle h, \omega \rangle} - 1|^2 d^2\omega \\ &= \frac{1}{4\pi^2} \iint_{\mathcal{B}_\infty(\varepsilon/2)} |\widehat{f_0}(\omega)|^2 (2 - 2\cos\langle h, \omega \rangle) d^2\omega. \end{aligned}$$

The integral is computed on a compact domain because, according to Lemma 1, $\text{supp } \widehat{f_0} \subset \mathcal{B}_\infty(\varepsilon/2)$. Now, we use the Cauchy-Schwarz inequality to compute:

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

By hypothesis on h , $\frac{\varepsilon}{2} \|h\|_1 \leq \frac{\pi}{2}$, and thus

$$\|\mathcal{T}_h f_0 - f_0\|_{L^2}^2 \leq \left(2 - 2\cos\frac{\varepsilon\|h\|_1}{2}\right) \|f_0\|_{L^2}^2.$$

Finally, since $\cos t \geq 1 - \frac{t^2}{2}$, we get (7). \blacksquare

APPENDIX C PROOF OF LEMMA 2

Proof: Expression (9) is obtained by adapting (3.2) and (3.3) in [48] to the 2D case. Then, by combining (9) with the Plancherel formula, we get

$$\begin{aligned} \|g\|_{L^2}^2 &= \frac{1}{4\pi^2} \|\widehat{g}\|_{L^2}^2 \\ &= \frac{1}{4\pi^2} \iint_{\mathcal{B}_\infty(\pi/s)} |\widehat{g}(\omega)|^2 d^2\omega \\ &= \frac{1}{4\pi^2} \iint_{\mathcal{B}_\infty(\pi/s)} |s \widehat{Y}(s\omega)|^2 d^2\omega. \end{aligned}$$

The integral is performed on $\mathcal{B}_\infty(\pi/s)$ because $g \in \mathcal{U}_s$. Then, by applying the change of variable $\omega' \leftarrow s\omega$, we get

$$\begin{aligned} \|g\|_{L^2}^2 &= \frac{1}{4\pi^2} \iint_{\mathcal{B}_\infty(\pi)} |\widehat{Y}(\omega')|^2 d^2\omega' \\ &= \frac{1}{4\pi^2} \|\widehat{Y}\|_{L^2}^2 = \|Y\|_2^2, \end{aligned}$$

hence (10), which concludes the proof. \blacksquare

APPENDIX D
PROOF OF PROPOSITION 2

Proof: First, f_X and ψ_W are well defined because $X \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$ and $W \in l^2(\mathbb{Z}^2)$. By construction, $\psi_W \in \mathcal{U}_s$. Thus, (12) is a direct consequence of the Shannon-Whittaker sampling theorem [48, p. 61], coupled with the orthonormality of $\{\phi_n^{(s)}\}_{n \in \mathbb{Z}^2}$. Therefore, using (9) in Lemma 2, we get, for any $\omega \in \mathcal{B}_{\infty}(\pi/s)$, $\widehat{\psi}_W(\omega) = s\widehat{W}(s\omega)$. Since $\widehat{\psi}_W(\omega) = 0$ outside $\mathcal{B}_{\infty}(\pi/s)$, we deduce that $\psi_W \in \mathcal{V}(\xi/s, \kappa/s)$.

We now prove (13). For $\mathbf{n} \in \mathbb{Z}^2$, we compute:

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

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

$$\begin{aligned} (f_X * \bar{\psi}_W)(ms\mathbf{n}) &= \sum_{\mathbf{k} \in \mathbb{Z}^2} X[\mathbf{k}] \iint_{\mathbb{R}^2} \bar{\psi}(\mathbf{x}) \phi_{\mathbf{k}}^{(s)}(s(m\mathbf{n} - \mathbf{k}) - \mathbf{x}) d^2\mathbf{x} \\ &= \sum_{\mathbf{k} \in \mathbb{Z}^2} X[\mathbf{k}] (\bar{\psi} * \phi^{(s)})(s(m\mathbf{n} - \mathbf{k})). \end{aligned}$$

Since $\{\phi_n^{(s)}\}_{n \in \mathbb{Z}^2}$ is an orthonormal basis of \mathcal{U}_s , we can easily show that, for any $\mathbf{k}' \in \mathbb{Z}^2$, $\overline{W}[\mathbf{k}'] = \langle \bar{\psi}, \phi_{\mathbf{k}'}^{(s)} \rangle = (\bar{\psi} * \phi^{(s)})(s\mathbf{k}')$. Therefore we get

$$\begin{aligned} (f_X * \bar{\psi}_W)(ms\mathbf{n}) &= \sum_{\mathbf{k} \in \mathbb{Z}^2} X[\mathbf{k}] \overline{W}[m\mathbf{n} - \mathbf{k}] \\ &= (X * \overline{W})[m\mathbf{n}], \end{aligned}$$

hence the result. \blacksquare

APPENDIX E
PROOF OF LEMMA 3

Proof: Let $\mathbf{u} \in \mathbb{R}^2$. By definition of f_X , $f_{\mathcal{T}_u X}$ and $\mathcal{T}_u X$,

$$f_{\mathcal{T}_u X} = \sum_{\mathbf{n} \in \mathbb{Z}^2} s \mathcal{T}_{su} f_X(s\mathbf{n}) \phi_{\mathbf{n}}^{(s)}. \quad (93)$$

By construction, $f_X \in \mathcal{U}_s$. Therefore, $\mathcal{T}_{su} f_X \in \mathcal{U}_s$. Then, the Shannon-Whittaker theorem [48, p. 61], coupled with the orthonormality of $\{\phi_n^{(s)}\}_{n \in \mathbb{Z}^2}$, implies

$$s \mathcal{T}_{su} f_X(s\mathbf{n}) = \langle \mathcal{T}_{su} f_X, \phi_{\mathbf{n}}^{(s)} \rangle. \quad (94)$$

Finally, plugging (94) into (93) concludes the proof. \blacksquare

APPENDIX F
PROOF OF THEOREM 1

Proof: We consider

$$f_0 : \mathbf{x} \mapsto (f_X * \bar{\psi}_W)(\mathbf{x}) e^{i\langle \xi/s, \mathbf{x} \rangle}, \quad (95)$$

with f_X and ψ_W satisfying (11). Then,

$$|f_X * \bar{\psi}_W| = |f_0| \quad \text{and} \quad |\mathcal{T}_{su} f_X * \bar{\psi}_W| = |\mathcal{T}_{su} f_0|. \quad (96)$$

According to Proposition 2, $\psi_W \in \mathcal{V}(\xi/s, \kappa/s)$. Therefore, according to Lemma 1, $\text{supp } \widehat{f_0} \subset \mathcal{B}_{\infty}(\frac{\kappa}{2s})$. Moreover, by hypothesis, $\kappa \leq 2\pi/m$; thus, $\mathcal{B}_{\infty}(\frac{\kappa}{2s}) \subset \mathcal{B}_{\infty}(\frac{\pi}{ms})$. Therefore, $f_0 \in \mathcal{U}_{s'}$, where we have denoted $s' := ms$.

According to (13) (Proposition 2), (18) (Corollary 1) and (96), we get

$$\mathcal{C}_m X[\mathbf{n}] = |f_0(s'\mathbf{n})|; \quad (97)$$

$$\mathcal{C}_m(\mathcal{T}_u X)[\mathbf{n}] = |(\mathcal{T}_{su} f_0)(s'\mathbf{n})|. \quad (98)$$

Then, using the reverse triangular inequality on (97) and (98),

$$\begin{aligned} \|\mathcal{C}_m(\mathcal{T}_u X) - \mathcal{C}_m X\|_2^2 &\leq \sum_{\mathbf{n} \in \mathbb{Z}^2} |\mathcal{T}_{su} f_0(s'\mathbf{n}) - f_0(s'\mathbf{n})|^2 \\ &= \sum_{\mathbf{n} \in \mathbb{Z}^2} |g(s'\mathbf{n})|^2 = \frac{1}{s'^2} \|Y\|_2^2, \end{aligned}$$

where we have denoted, for any $\mathbf{n} \in \mathbb{Z}^2$, $g := \mathcal{T}_{su} f_0 - f_0$ and $Y[\mathbf{n}] := s'g(s'\mathbf{n})$. We have $g \in \mathcal{U}_{s'}$ since $f_0 \in \mathcal{U}_{s'}$. Then, according to (10) in Lemma 2, $\|Y\|_2 = \|g\|_{L^2}$. Therefore,

$$\|\mathcal{C}_m(\mathcal{T}_u X) - \mathcal{C}_m X\|_2^2 \leq \frac{1}{s'^2} \|g\|_{L^2}^2 = \frac{1}{s'^2} \|\mathcal{T}_{su} f_0 - f_0\|_{L^2}^2.$$

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

$$\|\mathcal{C}_m(\mathcal{T}_u X) - \mathcal{C}_m X\|_2^2 \leq \frac{\alpha(\kappa\mathbf{u})^2}{s'^2} \|f_0\|_{L^2}^2. \quad (99)$$

Besides, according again to Lemma 2, $\|f_0\|_{L^2}^2 = \|X_0\|_2^2$, where $X_0[\mathbf{n}] := s'f_0(s'\mathbf{n})$ for any $\mathbf{n} \in \mathbb{Z}^2$. Therefore, according to (97),

$$\|\mathcal{C}_m X\|_2 = \frac{1}{s'} \|X_0\|_2 = \frac{1}{s'} \|f_0\|_{L^2}. \quad (100)$$

Finally, plugging (100) into (99) completes the proof. \blacksquare

APPENDIX G
PROOF OF PROPOSITION 3

Proof: Let $X \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$ and $s > 0$. We consider $f_0 \in L^2(\mathbb{R}^2)$ as the ‘‘low-frequency’’ function satisfying (95). Again, we introduce $s' := ms$ and $X_0 \in l^2(\mathbb{Z}^2)$ such that $X_0[\mathbf{n}] := s'f_0(s'\mathbf{n})$. Moreover, for any $Y \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$, we denote:

$$f_Y^{(s')} := \sum_{\mathbf{n} \in \mathbb{Z}^2} Y[\mathbf{n}] \phi_{\mathbf{n}}^{(s')}. \quad (101)$$

On the one hand, in the proof of Theorem 1, we already got (100). On the other hand, (98) can be rewritten

$$\mathcal{C}_m(\mathcal{T}_u X)[\mathbf{n}] = |\mathcal{T}_{s'\mathbf{u}'} f_0(s'\mathbf{n})|, \quad (102)$$

with $\mathbf{u}' := \mathbf{u}/m$. Besides, according to the proof of Theorem 1, $f_0 \in \mathcal{U}_{s'}$. Thus, by definition of X_0 , the Shannon-Whittaker theorem [48, p. 61] implies that $f_0 = f_{X_0}^{(s')}$ such as defined in (101). Then, using Lemma 3 with $X \leftarrow X_0$, $\mathbf{u} \leftarrow \mathbf{u}'$ and $s \leftarrow s'$, we get

$$f_{\mathcal{T}_{\mathbf{u}'} X_0}^{(s')} = \mathcal{T}_{s'\mathbf{u}'} f_{X_0}^{(s')} = \mathcal{T}_{s'\mathbf{u}'} f_0. \quad (103)$$

Then, using (12) in Proposition 2 with $X \leftarrow \mathcal{T}_{\mathbf{u}'} X_0$ and $s \leftarrow s'$, and inserting (103) into the result yields

$$\mathcal{T}_{\mathbf{u}'} X_0[\mathbf{n}] = s' f_{\mathcal{T}_{\mathbf{u}'} X_0}^{(s')}(s'\mathbf{n}) = s' \mathcal{T}_{s'\mathbf{u}'} f_0(s'\mathbf{n}). \quad (104)$$

Therefore, (102) and (104) imply

$$\|\mathcal{C}_m(\mathcal{T}_u X)\|_2 = \frac{1}{s'} \|\mathcal{T}_{u'} X_0\|_2. \quad (105)$$

Moreover, since $f_0 \in \mathcal{U}_{s'}$, and according to (104), we can use Lemma 2 with $s \leftarrow s'$, $g \leftarrow \mathcal{T}_{s'u'} f_0$ and $Y \leftarrow \mathcal{T}_{u'} X_0$. We get

$$\|\mathcal{T}_{u'} X_0\|_2 = \|\mathcal{T}_{s'u'} f_0\|_{L^2} = \|f_0\|_{L^2}. \quad (106)$$

Finally, (100), (105) and (106) imply $\|\mathcal{C}_m(\mathcal{T}_u X)\|_2 = \|\mathcal{C}_m X\|_2$, which concludes the proof. \blacksquare

APPENDIX H

JUSTIFICATION FOR HYPOTHESES 1 AND 2

Given $n \in \mathbb{N}^*$, we define n -th order stationarity of a given stochastic process F as in [55, p. 152]: for any $n' \in \{1..n\}$, $(\mathbf{x}_1, \dots, \mathbf{x}_{n'}) \in (\mathbb{R}^2)^{n'}$ and $\mathbf{h} \in \mathbb{R}^2$, the joint distribution of $(F(\mathbf{x}_1), \dots, F(\mathbf{x}_{n'}))$ is identical to the one of $(F(\mathbf{x}_1 + \mathbf{h}), \dots, F(\mathbf{x}_{n'} + \mathbf{h}))$. Besides, *strict-sense stationarity* is defined as n -th order stationarity for any $n \in \mathbb{N}^*$.

We recall that $\nu := \xi/s$. We then state the following results.

Proposition 8: We assume that F_X is first-order stationary. If, for any $\mathbf{x} \in \mathbb{R}^2$ and any $\mathbf{h} \in \mathcal{B}_2(2\pi/\|\nu\|_2)$,

$$(\mathcal{T}_h F_X * \bar{\psi})(\mathbf{x}) = e^{i\langle \nu, \mathbf{h} \rangle} (F_X * \bar{\psi})(\mathbf{x}), \quad (107)$$

then Hypothesis 1 is satisfied.

Proof: We show that the probability measure of $Z_X(\mathbf{x})$ is invariant with respect to phase shift. In other words, we show that, for any measurable set $\mathfrak{A} \subset \mathbb{S}^1$,

$$\forall \omega \in [0, 2\pi], \mu(\mathfrak{A}) = \mu(e^{i\omega} \mathfrak{A}), \quad (108)$$

where we have denoted

$$\mu : \mathfrak{A} \mapsto \mathbb{P}\{Z_X(\mathbf{x}) \in \mathfrak{A}\}.$$

Let $\mathbf{h} \in \mathcal{B}_2(2\pi/\|\nu\|_2)$. According to (107),

$$Z_X(\mathbf{x}) \in \mathfrak{A} \iff \mathcal{T}_h Z_X(\mathbf{x}) \in e^{i\langle \nu, \mathbf{h} \rangle} \mathfrak{A}.$$

Therefore,

$$\mathbb{P}\{Z_X(\mathbf{x}) \in \mathfrak{A}\} = \mathbb{P}\{\mathcal{T}_h Z_X(\mathbf{x}) \in e^{i\langle \nu, \mathbf{h} \rangle} \mathfrak{A}\}.$$

Since F_X is first-order stationary, $Z_X(\mathbf{x})$ and $\mathcal{T}_h Z_X(\mathbf{x})$ have the same probability distribution. Thus we get

$$\mathbb{P}\{Z_X(\mathbf{x}) \in \mathfrak{A}\} = \mathbb{P}\{Z_X(\mathbf{x}) \in e^{i\langle \nu, \mathbf{h} \rangle} \mathfrak{A}\}.$$

Let $\omega \in [0, 2\pi]$. Considering $\mathbf{h} := \omega \nu / \|\nu\|_2^2$, we have $\mathbf{h} \in \mathcal{B}_2(2\pi/\|\nu\|_2)$ and $\langle \nu, \mathbf{h} \rangle = \omega$. Therefore,

$$\forall \omega \in [0, 2\pi], \mathbb{P}\{Z_X(\mathbf{x}) \in \mathfrak{A}\} = \mathbb{P}\{Z_X(\mathbf{x}) \in e^{i\omega} \mathfrak{A}\},$$

which yields (108).

Any probability measure defined on \mathbb{S}^1 is a Radon measure. Therefore, according to Haar's theorem [56], there exists a unique probability measure on \mathbb{S}^1 satisfying (108). Since the uniform probability measure is also invariant to phase shift, we deduce that $Z_X(\mathbf{x})$ is uniformly distributed on \mathbb{S}^1 . \blacksquare

Proposition 9: We assume the conditions of Proposition 8 are met. If, moreover, F_X is strict-sense stationary, then Hypothesis 2 is satisfied.

Proof: Let $n \in \mathbb{N}^*$ and $\mathbf{x}, \mathbf{y}_1, \dots, \mathbf{y}_n \in \mathbb{R}^2$. To alleviate notations, we consider the random vector $\mathbf{M} = (M_X(\mathbf{y}_i))_{i \in \{1..n\}}$ with outcomes in \mathbb{R}_+^n .

The proof is organized as follows. Using a similar reasoning as Proposition 8, we show that Z_X follows a uniform conditional probability distribution given \mathbf{M} . Since we already know that Z_X follows a uniform (unconditional) distribution, we deduce that Z_X and \mathbf{M} are independent.

Let $\mathfrak{A} \subset \mathbb{S}^1$ and $\mathfrak{S} := (\mathfrak{S}_i)_{i \in \{1..n\}} \subset \mathbb{R}_+^n$ denote measurable sets. Let $\mathbf{h} \in \mathcal{B}_2(2\pi/\|\nu\|_2)$. According to (107),

$$Z_X(\mathbf{x}) \in \mathfrak{A} \iff \mathcal{T}_h Z_X(\mathbf{x}) \in e^{i\langle \nu, \mathbf{h} \rangle} \mathfrak{A};$$

$$M_X(\mathbf{y}_i) \in \mathfrak{S}_i \iff \mathcal{T}_h M_X(\mathbf{y}_i) \in \mathfrak{S}_i \quad \forall i \in \{1..n\}.$$

Therefore,

$$\mathbb{P}\{(Z_X(\mathbf{x}) \in \mathfrak{A}) \& (\mathbf{M} \in \mathfrak{S})\} =$$

$$\mathbb{P}\{(\mathcal{T}_h Z_X(\mathbf{x}) \in e^{i\langle \nu, \mathbf{h} \rangle} \mathfrak{A}) \& (\mathcal{T}_h \mathbf{M} \in \mathfrak{S})\}.$$

Since F_X is strict-sense stationary, the joint probability density of $\mathcal{T}_h Z_X(\mathbf{x}), \mathcal{T}_h M_X(\mathbf{y}_1), \dots, \mathcal{T}_h M_X(\mathbf{y}_n)$ is identical to the one of $Z_X(\mathbf{x}), M_X(\mathbf{y}_1), \dots, M_X(\mathbf{y}_n)$. Therefore we get

$$\mathbb{P}\{(Z_X(\mathbf{x}) \in \mathfrak{A}) \& (\mathbf{M} \in \mathfrak{S})\} =$$

$$\mathbb{P}\{(Z_X(\mathbf{x}) \in e^{i\langle \nu, \mathbf{h} \rangle} \mathfrak{A}) \& (\mathbf{M} \in \mathfrak{S})\}.$$

We assume that $\mathbb{P}(\mathbf{M} \in \mathfrak{S}) > 0$. According to the above expression, and similarly to the proof of Proposition 8, we get,

$$\forall \omega \in [0, 2\pi], \mathbb{P}\{Z_X(\mathbf{x}) \in \mathfrak{A} \mid \mathbf{M} \in \mathfrak{S}\} =$$

$$\mathbb{P}\{Z_X(\mathbf{x}) \in e^{i\omega} \mathfrak{A} \mid \mathbf{M} \in \mathfrak{S}\}.$$

Then, the above conditional probability measure satisfies phase shift invariance (108). Therefore, as in the proof of Proposition 8, Haar's theorem implies that $Z_X(\mathbf{x})$ follows a uniform conditional distribution given $\mathbf{M} \in \mathfrak{S}$.

Moreover, strict-sense implies first-order stationarity, and thus, according to Proposition 8, $Z_X(\mathbf{x})$ follows a uniform (unconditional) distribution. Therefore we get, for any measurable sets $\mathfrak{A} \subset \mathbb{S}^1$ and $\mathfrak{S} \subset \mathbb{R}_+^n$ such that $\mathbb{P}(\mathbf{M} \in \mathfrak{S}) > 0$,

$$\mathbb{P}\{Z_X \in \mathfrak{A} \mid \mathbf{M} \in \mathfrak{S}\} = \mathbb{P}(Z_X \in \mathfrak{A}),$$

which proves independence between $Z_X(\mathbf{x})$ and \mathbf{M} . \blacksquare

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 $L_{\mathbb{R}}^2(\mathbb{R}^2)$. 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 natural images can be found in [57]. 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 $\|\mathbf{h}\|_2 \ll sM$, where, as a reminder, M denotes the support size of input images.¹²

¹²We refer readers to [58] for a related notion of local stationarity.

As it turns out, the proofs of Propositions 8 and 9 only requires shifts with $\|\mathbf{h}\|_2 \leq 2\pi/\|\boldsymbol{\nu}\|_2$. Then, according to (53), $\|\mathbf{h}\|_2 \ll sM$, and the stationarity hypothesis holds.

Finally, to justify (107), we consider $\varphi_{\mathbb{W}} : \mathbf{x} \mapsto \psi_{\mathbb{W}}(\mathbf{x})e^{-i\langle \boldsymbol{\nu}, \mathbf{x} \rangle}$. Similarly to Lemma 1, we can show that $\varphi_{\mathbb{W}}$ is a low-pass filter, with $\text{supp } \widehat{\varphi}_{\mathbb{W}} \subset \mathcal{B}_{\infty}(\varepsilon/2)$. For all $\mathbf{h} \in \mathbb{R}^2$ such that $\|\mathbf{h}\|_2 \leq 2\pi/\|\boldsymbol{\nu}\|_2$, we have

$$\begin{aligned} & (\mathcal{T}_{\mathbf{h}}F_X * \overline{\psi})(\mathbf{x}) \\ &= \iint_{\mathbb{R}^2} \mathcal{T}_{\mathbf{h}}F_X(\mathbf{x} - \mathbf{y}) \overline{\varphi}(\mathbf{y}) e^{-i\langle \boldsymbol{\nu}, \mathbf{y} \rangle} d^2\mathbf{y} \\ &= e^{i\langle \boldsymbol{\nu}, \mathbf{h} \rangle} \iint_{\mathbb{R}^2} F_X(\mathbf{x} - \mathbf{y}') \overline{\varphi}(\mathbf{y}' - \mathbf{h}) e^{-i\langle \boldsymbol{\nu}, \mathbf{y}' \rangle} d^2\mathbf{y}'. \end{aligned}$$

Since $\text{supp } \widehat{\varphi}_{\mathbb{W}} \subset \mathcal{B}_{\infty}(\frac{\kappa}{2s})$, we can define a ‘‘minimal wavelength’’ $\lambda_{\varphi_{\mathbb{W}}} := 2\pi s/\kappa$. Then, if $\|\mathbf{h}\|_2 \ll \lambda_{\varphi_{\mathbb{W}}}$, we can approximate $\overline{\varphi}(\mathbf{y}' - \mathbf{h}) \approx \overline{\varphi}(\mathbf{y}')$. This sufficient condition is actually met, because $\|\mathbf{h}\|_2 \leq 2\pi/\|\boldsymbol{\nu}\|_2$ and, according to (54), $\|\boldsymbol{\nu}\|_2 \gg \kappa/s$. Therefore,

$$(\mathcal{T}_{\mathbf{h}}F_X * \overline{\psi})(\mathbf{x}) \approx e^{i\langle \boldsymbol{\nu}, \mathbf{h} \rangle} (F_X * \overline{\psi})(\mathbf{x}). \quad (109)$$

As a result, the conditions for Propositions 8 and 9 are approximately satisfied. We will therefore consider Hypotheses 1 and 2 as reasonable.

APPENDIX I

JUSTIFICATION FOR HYPOTHESIS 3

We consider, for any $l \in \mathcal{L}$ and any $k \in \{0 \dots K-1\}$, the value of $\mu \in \mathbb{R}$ minimizing $\|\mu \widetilde{V}_l - V_{lk}\|_2^2$, denoted by μ_{lk} . We then denote by $\delta_{lk} := \|\mu_{lk} \widetilde{V}_l - V_{lk}\|_2^2 / \|V_{lk}\|_2^2$ the relative quadratic error between V_{lk} and its projection on $\mathbb{R}\widetilde{V}_l$. We get

$$\delta_{lk} = 1 - \frac{\langle \widetilde{V}_l, V_{lk} \rangle^2}{\|\widetilde{V}_l\|_2^2 \cdot \|V_{lk}\|_2^2}. \quad (110)$$

Expression (67) holds if and only if $\delta_{lk} = 0$ for any $k \in \{0 \dots K-1\}$. In the case of AlexNet, when $l \in \mathcal{L}$, δ_{lk} do not exceed 10^{-1} , and its mean value is around 10^{-2} . Therefore $\delta_{lk} \ll 1$, and Hypothesis 3 can be considered as a reasonable assumption for any Gabor channel $l \in \mathcal{L}$. Similar observations can be drawn for ResNet.

APPENDIX J

PROOF OF PROPOSITION 4

Proof: We consider the Borel σ -algebra on \mathbb{S}^1 generated by $\{[z, z']_{\mathbb{S}^1} \mid z, z' \in \mathbb{S}^1\} \cup \{\mathbb{S}^1\}$, on which we have defined the angular measure ϑ such that $\vartheta(\mathbb{S}^1) := 2\pi$, and

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

For any $p \in \mathbb{N}^*$, we compute the p -th moment of $G_{\max X}(\mathbf{x})$ defined in (51). By considering

$$\begin{aligned} h_{\max} : \mathbb{S}^1 &\rightarrow [-1, 1] \\ z &\mapsto \max_{\|\mathbf{k}\|_{\infty} \leq q} \text{Re}(z^* z_{\mathbf{k}}), \end{aligned}$$

we get $G_{\max X}(\mathbf{x}) = h_{\max}(Z_X(\mathbf{x}))$. According to Hypothesis 1, $Z_X(\mathbf{x})$ follows a uniform distribution on \mathbb{S}^1 . Therefore,

$$\mathbb{E}[G_{\max X}(\mathbf{x})^p] = \frac{1}{2\pi} \int_{\mathbb{S}^1} h_{\max}(z)^p d\vartheta(z),$$

which proves that $\mathbb{E}[G_{\max X}(\mathbf{x})^p]$ does not depend on \mathbf{x} . Let us split the unit circle \mathbb{S}^1 into the arcs $\mathfrak{A}_0^{(q)}, \dots, \mathfrak{A}_{N_q-1}^{(q)}$ such as defined in (47):

$$\mathbb{E}[G_{\max X}(\mathbf{x})^p] = \frac{1}{2\pi} \sum_{i=0}^{N_q-1} \int_{\mathfrak{A}_i^{(q)}} h_{\max}(z)^p d\vartheta(z). \quad (111)$$

Let $i \in \{0 \dots N_q-1\}$. We can show that

$$\forall z \in \mathfrak{A}_i^{(q)}, h_{\max}(z) = \max\left(\text{Re}(z^* z_i^{(q)}), \text{Re}(z^* z_{i+1}^{(q)})\right).$$

Therefore, h_{\max} is symmetric with respect to the center value of $\mathfrak{A}_i^{(q)}$, denoted by $\bar{z}_i^{(q)}$, where h_{\max} reaches its minimum. We denote by $\overline{\mathfrak{A}}_i^{(q)} := [z_i^{(q)}, \bar{z}_i^{(q)}]_{\mathbb{S}^1}$ the first half of arc $\mathfrak{A}_i^{(q)}$. Then,

$$\forall z \in \overline{\mathfrak{A}}_i^{(q)}, h_{\max}(z) = \text{Re}(z^* z_i^{(q)}).$$

As a consequence, using symmetry, we get

$$\begin{aligned} \int_{\mathfrak{A}_i^{(q)}} h_{\max}(z)^p d\vartheta(z) &= 2 \int_{\overline{\mathfrak{A}}_i^{(q)}} h_{\max}(z)^p d\vartheta(z) \\ &= 2 \int_{\overline{\mathfrak{A}}_i^{(q)}} \text{Re}(z^* z_i^{(q)})^p d\vartheta(z). \end{aligned}$$

By using the change of variable formula [59, p. 81] with $z \leftarrow e^{i\theta}$, we get

$$\int_{\mathfrak{A}_i^{(q)}} h_{\max}(z)^p d\vartheta(z) = 2 \int_{\bar{\theta}_i^{(q)}}^{\theta_i^{(q)}} \cos^p(\theta - \theta_i^{(q)}) d\theta,$$

where $\bar{\theta}_i^{(q)} := (\theta_i^{(q)} + \theta_{i+1}^{(q)})/2$ denotes the argument of $\bar{z}_i^{(q)}$. Then, the change of variable $\omega \leftarrow \theta - \theta_i^{(q)}$ yields

$$\int_{\mathfrak{A}_i^{(q)}} h_{\max}(z)^p d\vartheta(z) = 2 \int_0^{\omega_i^{(q)}/2} \cos^p \omega d\omega. \quad (112)$$

Now, we insert (112) into (111), and compute $\mathbb{E}[G_{\max X}(\mathbf{x})^p]$ for $p=1$ and $p=2$. We get

$$\mathbb{E}[G_{\max X}(\mathbf{x})] = \frac{1}{\pi} \sum_{i=0}^{N_q-1} \sin \frac{\omega_i^{(q)}}{2};$$

$$\mathbb{E}[G_{\max X}(\mathbf{x})^2] = \frac{1}{2} + \frac{1}{4\pi} \sum_{i=0}^{N_q-1} \sin \omega_i^{(q)}.$$

We recall that $Q_X := 1 - G_{\max X}$. By linearity of \mathbb{E} , we get

$$\mathbb{E}[Q_X(\mathbf{x})^2] := \frac{3}{2} + \frac{1}{4\pi} \sum_{i=0}^{N_q-1} \left(\sin \omega_i^{(q)} - 8 \sin \frac{\omega_i^{(q)}}{2} \right).$$

Using the notation introduced in (48) concludes the proof. ■

APPENDIX K

PROOF OF PROPOSITION 5

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

$$\widetilde{X}_X^{(n)} := \sqrt{\sum_{\|\mathbf{k}\|_{\infty} \leq n} M_X(\mathbf{x}_{\mathbf{k}})^2}.$$

According to Hypothesis 2, $Z_X(\mathbf{x})$ is jointly independent of $M_X(\mathbf{x}_k)$ for $\mathbf{k} \in \{-n \dots n\}^2$. Therefore, by composition, $Z_X(\mathbf{x})$ is also independent of $\tilde{S}_X^{(n)}$. Moreover, according to (52) and (60), $\tilde{S}_X^{(n)}$ converges almost surely towards \tilde{S}_X , which proves independence between $Z_X(\mathbf{x})$ and \tilde{S}_X .

Now, we prove conditional independence between $Z_X(\mathbf{x})$ and $M_X(\mathbf{x})$ given \tilde{S}_X . According to Hypothesis 2,

$$(M_X(\mathbf{x}), \tilde{S}_X^{(n)}) \perp\!\!\!\perp Z_X(\mathbf{x}),$$

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

$$Z_X(\mathbf{x}) \perp\!\!\!\perp M_X(\mathbf{x}) \mid \tilde{S}_X^{(n)}.$$

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

APPENDIX L PROOF OF THEOREM 2

Proof: We consider $\mathbf{n} \in \mathbb{Z}^2$. By construction, $Q_X(\mathbf{x}_n) := 1 - G_{\max_X}(\mathbf{x}_n)$ only depends on $Z_X(\mathbf{x}_n)$. Therefore, according to Proposition 5, we have

$$Q_X(\mathbf{x}_n) \perp\!\!\!\perp M_X(\mathbf{x}_n) \mid \tilde{S}_X^2 \quad \text{and} \quad Q_X(\mathbf{x}_n) \perp\!\!\!\perp \tilde{S}_X^2. \quad (113)$$

Besides, we introduce $\tilde{\Delta}_X := \|\delta_{m,q}X\|_2$, where $\delta_{m,q}X$ is defined in (56). Then, using the linearity of \mathbb{E} , and according to (52), we get, for any $\sigma > 0$,

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

According to (113), we get

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

Under Hypothesis 1, we use (58) in Proposition 4, which yields $\mathbb{E}[Q_X(\mathbf{x}_n)^2] = \gamma_q(m\xi)^2$. Therefore, using again the linearity of \mathbb{E} , and according to (52) and (60),

$$\begin{aligned} \mathbb{E}[\tilde{\Delta}_X^2 \mid \tilde{S}_X^2 = \sigma] &= \mathbb{E}[\tilde{S}_X^2 \mid \tilde{S}_X^2 = \sigma] \cdot \gamma_q(m\xi)^2 \\ &= \sigma \cdot \gamma_q(m\xi)^2. \end{aligned}$$

Besides, we can reformulate \tilde{Q}_X such as defined in (57): $\tilde{Q}_X = \tilde{\Delta}_X / \tilde{S}_X$. Therefore we get

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

Then, law of total expectation states that

$$\mathbb{E}[\tilde{Q}_X^2] = \mathbb{E}[\mathbb{E}[\tilde{Q}_X^2 \mid \tilde{S}_X^2]] = \gamma_q(m\xi)^2.$$

Finally, according to Conjecture 1, we have:

$$\tilde{P}_X \leq (1 + \beta_q(m\kappa)) \tilde{Q}_X,$$

which yields (61). ■

APPENDIX M PROOF OF THEOREM 3

Proof: Using the triangular inequality, we compute

$$\begin{aligned} \|\mathcal{R}_{m,q}(\mathcal{T}_u X) - \mathcal{R}_{m,q}X\|_2 &\leq \|\mathcal{C}_{2m}(\mathcal{T}_u X)\|_2 \tilde{P}_{\mathcal{T}_u X} + \|\mathcal{C}_{2m}X\|_2 \tilde{P}_X \\ &\quad + \|\mathcal{C}_{2m}(\mathcal{T}_u X) - \mathcal{C}_{2m}X\|_2, \end{aligned}$$

where \tilde{P}_X and $\tilde{P}_{\mathcal{T}_u X}$ are defined in (55). Since, by hypothesis, $\kappa \leq \pi/m$, expression (21) in Proposition 3 states that

$$\|\mathcal{C}_{2m}(\mathcal{T}_u X)\|_2 = \|\mathcal{C}_{2m}X\|_2.$$

Moreover, according to (63), we can apply (20) in Theorem 1 on the third term of the above expression. We get

$$\begin{aligned} \|\mathcal{R}_{m,q}(\mathcal{T}_u X) - \mathcal{R}_{m,q}X\|_2 &\leq \left[\tilde{P}_{\mathcal{T}_u X} + \tilde{P}_X + \alpha(\kappa u) \right] \|\mathcal{C}_{2m}X\|_2. \end{aligned}$$

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

$$\mathbb{E}[\tilde{R}_{X,u}] \leq \mathbb{E}[\tilde{P}_{\mathcal{T}_u X}] + \mathbb{E}[\tilde{P}_X] + \alpha(\kappa u). \quad (114)$$

For any stochastic process X' satisfying Hypotheses 1 and 2, expression (61) in Theorem 2 and Jensen's inequality yield:

$$\mathbb{E}[\tilde{P}_{X'}] \leq (1 + \beta_q(m\kappa)) \gamma_q(m\xi). \quad (115)$$

Since Hypotheses 1 & 2 are satisfied for Z_X and M_X , Lemma 3 implies that they are also true for $Z_{\mathcal{T}_u X}$ and $M_{\mathcal{T}_u X}$. Therefore, (115) is valid for $X' \leftarrow X$ and $X' \leftarrow \mathcal{T}_u X$, and plugging this expression into (114) concludes the proof. ■

APPENDIX N PROOF OF PROPOSITION 6

Proof: This proposition is a simple reformulation of the well-known result that two successive convolutions can be written as another convolution with a wider kernel. We introduce the upsampling operator: $(X \uparrow m)[\mathbf{n}] := X[\mathbf{n}/m]$ if $\mathbf{n}/m \in \mathbb{Z}^2$, and 0 otherwise. We also consider the ‘‘identity’’ filter $I \in l^2(\mathbb{Z}^2)$ such that $I[0] = 1$ and $I[\mathbf{n}] = 0$ otherwise.

First, for any $U, V \in l^2(\mathbb{Z}^2)$ and any $s, t \in \mathbb{N}^*$, we have $((U \downarrow s) * V) \downarrow t = (U * (V \uparrow s)) \downarrow (st)$. Then, a simple reasoning by induction yields the result, with

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

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

APPENDIX O PROOF OF PROPOSITION 7

Proof: For any $k \in \{0 \dots 3\}$, Proposition 6 guarantees the existence of $V_{k,l}^{(J)} \in l^2_{\mathbb{R}}(\mathbb{Z}^2)$ such that

$$X_{k,l}^{(J)} = (X * \overline{V_{k,l}^{(J)}}) \downarrow 2^J. \quad (116)$$

Then, the result is obtained by plugging (116) into (82) for $k \in \{0 \dots 3\}$, and by denoting

$$\begin{pmatrix} W_l^{(J)} \\ W_{4^J+l}^{(J)} \end{pmatrix} := \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} V_{0,l}^{(J)} \\ V_{3,l}^{(J)} \end{pmatrix} + i \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} V_{2,l}^{(J)} \\ V_{1,l}^{(J)} \end{pmatrix}.$$

■

■

REFERENCES

- [1] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [2] M. Vetterli, "Wavelets, approximation, and compression," *IEEE Signal Processing Magazine*, vol. 18, no. 5, pp. 59–73, Sep. 2001.
- [3] A. Laine and J. Fan, "Texture classification by wavelet packet signatures," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 15, no. 11, pp. 1186–1191, Nov. 1993.
- [4] S. Pittner and S. V. Kamarthi, "Feature extraction from wavelet coefficients for pattern recognition tasks," *IEEE Trans. Pattern analysis and machine intelligence*, vol. 21, no. 1, pp. 83–88, Jan. 1999.
- [5] G. G. Yen, "Wavelet packet feature extraction for vibration monitoring," *IEEE Trans. Industrial Electronics*, vol. 47, no. 3, pp. 650–667, Jun. 2000.
- [6] K. Huang and S. Aviyente, "Wavelet feature selection for image classification," *IEEE Trans. Image Processing*, vol. 17, no. 9, pp. 1709–1720, Aug. 2008.
- [7] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2323, Nov. 1998.
- [8] T. Wiatowski and H. Bölcskei, "A mathematical theory of deep convolutional neural networks for feature extraction," *IEEE Trans. Information Theory*, vol. 64, no. 3, pp. 1845–1866, Mar. 2018.
- [9] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *ICLR*, 2015.
- [10] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," *Communications of the ACM*, vol. 60, no. 6, pp. 84–90, May 2017.
- [11] J. Bruna and S. Mallat, "Invariant scattering convolution networks," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 35, no. 8, pp. 1872–1886, May 2013.
- [12] B. Liao and F. Peng, "Rotation-invariant texture features extraction using dual-tree complex wavelet transform," in *Intl. Conf. Information, Networking and Automation (ICINA)*, 2010.
- [13] L. Sifre and S. Mallat, "Rotation, scaling and deformation invariant scattering for texture discrimination," in *CVPR*, 2013.
- [14] A. Bietti and J. Mairal, "Invariance and stability of deep convolutional representations," in *NeurIPS*, 2017.
- [15] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, "How transferable are features in deep neural networks?" in *NeurIPS*, 2014.
- [16] M. Rai and P. Rivas, "A review of convolutional neural networks and Gabor filters in object recognition," in *Intl. Conf. Computational Science and Computational Intelligence (CSCI)*, 2020.
- [17] S.-Y. Chang and N. Morgan, "Robust CNN-based speech recognition with Gabor filter kernels," in *INTERSPEECH*, 2014.
- [18] S. Fujieda, K. Takayama, and T. Hachisuka, "Wavelet convolutional neural networks for texture classification," *arXiv:1707.07394*, Jul. 2017.
- [19] S. S. Sarwar, P. Panda, and K. Roy, "Gabor filter assisted energy efficient fast learning convolutional neural networks," in *IEEE/ACM International Symposium on Low Power Electronics and Design*, 2017.
- [20] S. Luan, C. Chen, B. Zhang, J. Han, and J. Liu, "Gabor convolutional networks," *IEEE Trans. Image Processing*, vol. 27, no. 9, pp. 4357–4366, May 2018.
- [21] P. Liu, H. Zhang, W. Lian, and W. Zuo, "Multi-level wavelet convolutional neural networks," *IEEE Access*, vol. 7, pp. 74 973–74 985, Jun. 2019.
- [22] M. Ulicny, V. A. Krylov, and R. Dahyot, "Harmonic networks for image classification," in *BMVC*, 2019.
- [23] A. Azulay and Y. Weiss, "Why do deep convolutional networks generalize so poorly to small image transformations?" *JMLR*, vol. 20, no. 184, pp. 1–25, 2019.
- [24] R. Zhang, "Making convolutional networks shift-invariant again," in *ICML*, 2019.
- [25] C. Vasconcelos, H. Larochelle, V. Dumoulin, N. L. Roux, and R. Goroshin, "An effective anti-aliasing approach for residual networks," *arXiv:2011.10675*, Nov. 2020.
- [26] X. Zou, F. Xiao, Z. Yu, and Y. J. Lee, "Delving Deeper into Anti-aliasing in ConvNets," in *BMVC*, 2020.
- [27] A. Chaman and I. Dokmanic, "Truly shift-invariant convolutional neural networks," in *CVPR*, 2021.
- [28] I. Waldspurger, "Wavelet transform modulus : Phase retrieval and scattering," Doctoral Thesis, Ecole normale supérieure, Paris, 2015.
- [29] N. Kingsbury, "Complex wavelets for shift invariant analysis and filtering of signals," *Applied and computational harmonic analysis*, vol. 10, no. 3, pp. 234–253, May 2001.
- [30] I. Bayram and I. W. Selesnick, "On the dual-tree complex wavelet packet and M-band transforms," *IEEE Trans. Signal Processing*, vol. 56, no. 6, pp. 2298–2310, Jun. 2008.
- [31] H. Leterme, K. Polisano, V. Perrier, and K. Alahari, "Modélisation parcimonieuse de CNNs avec des paquets d'ondelettes dual-tree," in *ORASIS*, 2021.
- [32] N. Kingsbury and J. Magarey, "Wavelet transforms in image processing," in *Signal Analysis and Prediction*. Birkhäuser, 1998, pp. 27–46.
- [33] E. Oyallon and S. Mallat, "Deep roto-translation scattering for object classification," in *CVPR*, 2015.
- [34] A. Singh and N. Kingsbury, "Dual-Tree wavelet scattering network with parametric log transformation for object classification," in *ICASSP*, 2017.
- [35] D. Zou and G. Lerman, "Graph convolutional neural networks via scattering," *Applied and Computational Harmonic Analysis*, vol. 49, no. 3, pp. 1046–1074, Nov. 2020.
- [36] E. Oyallon, E. Belilovsky, and S. Zagoruyko, "Scaling the scattering transform: Deep hybrid networks," in *ICCV*, 2017.
- [37] E. Oyallon, E. Belilovsky, S. Zagoruyko, and M. Valko, "Compressing the input for CNNs with the first-order scattering transform," in *ECCV*, 2018.
- [38] J. Zarka, L. Thiry, T. Angles, and S. Mallat, "Deep network classification by scattering and homotopy dictionary learning," in *ICLR*, 2020.
- [39] J. Zarka, F. Guth, and S. Mallat, "Separation and concentration in deep networks," in *ICLR*, 2021.
- [40] F. Cotter and N. Kingsbury, "A learnable scatternet: Locally invariant convolutional layers," in *ICIP*, 2019.
- [41] S. Gauthier, B. Thérien, L. Alsène-Racicot, M. Chaudhary, I. Rish, E. Belilovsky, M. Eickenberg, and G. Wolf, "Parametric scattering networks," in *CVPR*, 2022.
- [42] S. Mallat, "Group invariant scattering," *Communications on Pure and Applied Mathematics*, vol. 65, no. 10, pp. 1331–1398, Jul. 2012.
- [43] —, "Understanding deep convolutional networks," *Philosophical Trans. of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 374, no. 2065, Apr. 2016.
- [44] W. Czaja and W. Li, "Analysis of time-frequency scattering transforms," *Applied and Computational Harmonic Analysis*, vol. 47, no. 1, pp. 149–171, Jul. 2019.
- [45] —, "Rotationally invariant time–frequency scattering transforms," *J. Fourier Analysis and Applications*, vol. 26, no. 1, p. 4, Jan. 2020.
- [46] A. Bietti and J. Mairal, "Group invariance, stability to deformations, and complexity of deep convolutional representations," *JMLR*, vol. 20, no. 1, pp. 876–924, 2019.
- [47] J. Havlicek, J. Havlicek, and A. Bovik, "The analytic image," in *ICIP*, 1997.
- [48] S. Mallat, *A Wavelet Tour of Signal Processing : The Sparse Way*. Academic Press, 2009.
- [49] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *ICML*, 2015.
- [50] N. Kingsbury, "Design of Q-shift complex wavelets for image processing using frequency domain energy minimization," in *ICIP*, 2003.
- [51] G. P. Nason and B. W. Silverman, "The stationary wavelet transform and some statistical applications," in *Wavelets and Statistics*, ser. Lecture Notes in Statistics. Springer, 1995, pp. 281–299.
- [52] Y. Meyer, "Principe d'incertitude, bases hilbertiennes et algèbres d'opérateurs," in *Séminaire Bourbaki*, vol. 662, 1985.
- [53] I. W. Selesnick, R. Baraniuk, and N. Kingsbury, "The dual-tree complex wavelet transform," *IEEE Signal Processing Magazine*, vol. 22, no. 6, pp. 123–151, Nov. 2005.
- [54] 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," *IJCV*, vol. 115, no. 3, pp. 211–252, Apr. 2015.
- [55] K. I. Park and M. Park, *Fundamentals of Probability and Stochastic Processes with Applications to Communications*. Springer, 2018.
- [56] P. R. Halmos, *Measure Theory*. Springer, 2013.
- [57] A. Torralba and A. Oliva, "Statistics of natural image categories," *Network*, vol. 14, no. 3, pp. 391–412, Jan. 2003.
- [58] M. Tytgert, J. Bruna, S. Chintala, Y. LeCun, S. Piantino, and A. Szlam, "A mathematical motivation for complex-valued convolutional networks," *Neural Computation*, vol. 28, no. 5, pp. 815–825, May 2016.
- [59] K. B. Athreya and S. N. Lahiri, *Measure Theory and Probability Theory*. Springer, 2006, vol. 19.