
BEYOND SPARSE DENOISING IN FRAMES: MINIMAX ESTIMATION WITH A SCATTERING TRANSFORM

Nathanaël Cuvelle-Magar*
Ecole normale supérieure
nmagar@di.ens.fr

Stéphane Mallat
Collège de France
Ecole normale supérieure
Flatiron Institute, Simons Foundation
stephane.mallat@ens.fr

ABSTRACT

A considerable amount of research in harmonic analysis has been devoted to non-linear estimators of signals contaminated by additive Gaussian noise. They are implemented by thresholding coefficients in a frame, which provide a sparse signal representation, or by minimising their ℓ^1 norm. However, sparse estimators in frames are not sufficiently rich to adapt to complex signal regularities. For cartoon images whose edges are piecewise C^α curves, wavelet, curvelet and Xlet frames are suboptimal if the Lipschitz exponent $\alpha \leq 2$ is an unknown parameter. Deep convolutional neural networks have recently obtained much better numerical results, which reach the minimax asymptotic bounds for all α . Wavelet scattering coefficients have been introduced as simplified convolutional neural network models. They are computed by transforming the modulus of wavelet coefficients with a second wavelet transform. We introduce a denoising estimator by jointly minimising and maximising the ℓ^1 norms of different subsets of scattering coefficients. We prove that these ℓ^1 norms capture different types of geometric image regularity. Numerical experiments show that this denoising estimator reaches the minimax asymptotic bound for cartoon images for all Lipschitz exponents $\alpha \leq 2$. We state this numerical result as a mathematical conjecture. It provides a different harmonic analysis approach to suppress noise from signals, and to specify the geometric regularity of functions. It also opens a mathematical bridge between harmonic analysis and denoising estimators with deep convolutional network.

Keywords Sparse denoising · minimax estimation · geometric regularity · wavelets

1 Introduction

Signals may be contaminated by additive Gaussian white noise during measurements or transmission. Such noise removal has thus been a central signal processing topic since the early work of Wiener on optimal linear denoisers [56]. Donoho and Johnstone [14] have proved that nearly optimal denoising estimators are obtained over signal classes which have sparse decomposition coefficients over a basis or a frame. The denoising estimator is computed by thresholding the noisy signal coefficients in this frame, or by minimising their ℓ^1 norm as explained in Section 2.1. In wavelet bases, it defines nearly minimax estimators over balls of Besov spaces [15], which correspond to functions having specific local isotropic regularity properties. Section 2.2 gives a brief review. For images having anisotropic directional regularities, a considerable effort has been devoted to build frames such as curvelets or shearlets, which define sparse representations. However, bases and frames have a limited flexibility to adapt to various geometric regularities. A major issue in harmonic analysis addressed in this paper is to develop new mathematical tools to capture sparsity properties and improve these estimators. The paper gives a mathematical conjecture sustained by numerical results and partial mathematical justifications to do so, after reviewing important harmonic analysis results.

Images with edges having some geometric regularity are important examples introduced by Tsybakov and Korostelev [30]. We shall consider an image model [45] of piecewise C^α (uniformly Lipschitz α) geometric images, which are C^α

*Corresponding author.

over disjoint image domains separated by regular edges. These edges are themselves regular C^α curves. Edge junctions define corners which are geometric singularities. Sparse representations are obtained with curvelets and shearlets if $\alpha = 2$ and if there is no corner, as explained in Section 2.3. However, there is no frame providing an optimally sparse representation for such images having corners, for a Lipschitz exponent $\alpha \leq 2$ which may vary. Indeed, it would require to adapt the geometry of frame vectors to the geometric regularity of edges in the image. If the edge regularity varies then we must choose the decomposition vectors in dictionaries of functions which are much larger than a frame. Optimally sparse representations have been obtained by selecting an orthonormal basis in a dictionary of bandlets [45, 18], whose geometry is adapted to each noisy image in terms of location, orientation and degree of regularity of edges. However, this approach becomes progressively too complex as the signal diversity increases, because the dictionary size must also increase.

More recently, Gaussian noise removal has gained considerable attention for AI data generation with score diffusion. New data are generated by transporting a Gaussian white noise with a stochastic differential equation which performs a progressive denoising, with a drift term computed from a denoiser implemented with a deep neural network [51, 24, 27]. Denoising capabilities of deep convolutional neural networks go well beyond the performance of previous estimators based on wavelets, curvelets, Xlets and bandlets but their mathematical properties are not understood.

A deep convolutional network computes a cascade of convolutional operators which have a narrow square support, and pointwise non-linearities. Simplified wavelet scattering models have been introduced in [41, 2] by replacing these convolutional operators with wavelet transforms and a modulus non-linearity. As in deep convolutional networks, all wavelet filters have a square support, as explained in Section 3.1. The main theorem of the paper in Section 3.2 proves that applying a second wavelet transform over the modulus of a first wavelet transform can still take advantage of the anisotropic geometric regularity of edges. This regularity appears through the fast decay of the ℓ^1 norm of a subset of scattering coefficients at fine scales. More surprisingly, the concentration of large wavelet coefficients along edge curves is also captured by the ℓ^1 norms of another subset of scattering coefficients, which become large and not small. Geometric image regularities are thus captured by both small and large ℓ^1 norms of scattering coefficients, which only involve wavelet filters.

A denoising algorithm is implemented by minimising a weighted average of these scattering ℓ^1 norms with positive and negative weights to insure a joint minimisation and maximisation of these norms. Numerical simulations in Section 3.3 demonstrate that this estimator reaches the minimax asymptotic rate over images having C^α edges for all values of $\alpha \leq 2$. All experiments are reproducible with the software available on <https://github.com/ScatteringDenoisingCode/SD>. Based on these numerical results, we state a mathematical conjecture on the minimax denoising rate of such denoising algorithm, which is not proved. This conjecture gives a new characterization of sparsity properties through minimisation and maximisation of different ℓ^1 norms. It does not require to use large size dictionaries with vectors whose support are adapted to the image geometry. It is a mathematical bridge towards deep convolutional network estimators.

2 Denoising geometrically regular images

This first section reviews the minimax optimality of denoising estimators, which relies on sparsity properties in harmonic analysis bases or frames. Section 2.1 reviews minimax estimations with regularised estimators promoting sparsity. Section 2.2 reviews sparse estimations in wavelet bases, which are suboptimal to remove noise from images having some form of geometric regularity. Section 2.3 reviews results obtained with curvelets and shearlets frames, bandlet dictionaries and deep convolutional neural networks.

2.1 Variational and sparse denoising estimations

Suppose that $f \in \mathbf{L}^2[0, 1]^2$ is contaminated by the addition of a Wiener measure of variance σ^2 (Gaussian white noise): $g = f + dB$, where B is a Brownian motion of variance σ^2 . We shall consider functions f in a set Λ having specific regularity properties. An estimator $\hat{f}(g)$ computes an approximation of f from g . It has a mean-squared error (MSE) $\mathbb{E}_g(\|\hat{f}(g) - f\|^2)$, where the expected value is computed over the Gaussian white noise distribution of g . The minimax error is the infimum which is achievable by optimizing the estimator $\hat{f}$

$$\epsilon_m(\sigma) = \inf_{\hat{f}} \sup_{f \in \Lambda} \mathbb{E}_g(\|\hat{f}(g) - f\|^2).$$

If f is a realization of a random process whose probability distribution is supported within Λ , then the minimum mean-squared error is calculated with an expected value over the distribution of f and g

$$\epsilon_{mms}(\sigma) = \inf_{\hat{f}} \mathbb{E}_{f,g}(\|\hat{f}(g) - f\|^2) \text{ and } \epsilon_{mms} \leq \epsilon_m.$$

Most work in harmonic analysis have been focused on the calculation of the minimax error ϵ_m which gives an upper bound on ϵ_{mms} overall distributions. Depending upon the properties of Λ , our goal is to define a denoising estimator $\hat{f}(g)$ whose maximum error over Λ

$$\epsilon_M(\sigma) = \sup_{f \in \Lambda} \mathbb{E}_g(\|\hat{f}(g) - f\|^2)$$

is close to the theoretical bounds computed over the minimax error ϵ_m . This maximum error is difficult to compute numerically. It is replaced in numerical experiments by a mean-squared error $\epsilon_{ms}(\sigma) = \mathbb{E}_{f,g}(\|\hat{f}(g) - f\|^2)$, for random processes whose typical realizations nearly reach the maximum error.

Regularised denoising estimators In numerical applications, the measurement device such as a camera gives a finite number d of coefficients which depend linearly on the input signal g . It thus characterises an orthogonal projection $P_V g$ of g in a vectorial subspace V of dimension d . The measurements project the signal as well as the noise. The projected noise $Z = P_V(dB)$ is a random Gaussian white noise vector of dimension d , so

$$P_V g = P_V f + Z \text{ where } Z \sim \mathcal{N}(0, \sigma^2 Id). \quad (1)$$

The problem is thus restricted to the space V of dimension d . This dimension varies depending on the noise level. Indeed, the resolution at which the original signal could be restored from its noisy counterpart is a function of σ .

We shall consider denoising estimators $\hat{f}$ of $P_V f$ from $P_V g$, which takes advantage of the prior information that f belongs to Λ . This may be captured by a regularity measure $U(f)$. Regularized estimators impose that $U(\hat{f})$ is bounded by $\sup_{f \in \Lambda} U(f)$. For example, a Tikhonov estimation [54, 55] assumes that Λ is a ball in a Sobolev space which is specified by a Sobolev norm $U(f)$. Given the regularity prior U , an estimation $\hat{f}(g) \in V$ of $P_V f$ is obtained as a solution of the variational problem

$$\hat{f}(g) = \operatorname{argmin}_{h \in V} \frac{1}{2} \|h - P_V g\|^2 + \sigma^2 U(h). \quad (2)$$

The error can be decomposed into an approximation and an estimation error

$$\|f - \hat{f}(g)\|^2 = \|f - P_V f\|^2 + \|P_V f - \hat{f}(g)\|^2. \quad (3)$$

In the following we will assume that the discretisation dimension d is large enough so that the leading term is the estimation error $\|P_V f - \hat{f}(g)\|^2$. When the prior set Λ is more complex than a ball in a Banach space, a major difficulty is to find $U(f)$ which defines a regularised estimator which is asymptotically minimax in Λ .

Sparsity with ℓ^1 norms in bases or frames An important body of research in applied harmonic analysis shows that nearly minimax estimators may be obtained by finding a basis $\{\psi_n\}_{n \leq d}$ of V , where the decomposition coefficients $Wf = \{\langle f, \psi_n \rangle\}_{n \leq d}$ are sufficiently sparse for all $f \in \Lambda$. This sparsity can be measured by an ℓ^1 norm of the vector Wf of decomposition coefficients

$$U(f) = \lambda \|Wf\|_1. \quad (4)$$

If the basis $\{\psi_n\}_{n \leq d}$ is orthonormal then one can verify [16, 1] that the solution $\hat{f}(g)$ of the variational problem (2) is computed by thresholding the coefficients of Wf with the soft-thresholding function

$$\rho_T(a) = \operatorname{sign}(a) \max(|a| - T, 0)$$

with $T = \lambda \sigma$. The estimation is calculated by applying the inverse W^T of W

$$\hat{f}(g) = W^T \rho_{\lambda \sigma}(Wg). \quad (5)$$

This thresholding can also be interpreted as a two layers neural network, whose weights are defined by W and W^T , with the point-wise discontinuity defined by the soft thresholding that can also be decomposed into a difference of two *ReLU*s. If the basis $\{\psi_n\}_{n \leq d}$ is not orthogonal in which case W in (4) is not unitary then $\hat{f}(g)$ is computed with an iterative proximal algorithm [9, 31]. Such algorithms can be efficiently approximated with a deep neural network [22, 36].

Asymptotic optimality Donoho and Johnstone [14, 13] have proved that thresholding estimators can achieve nearly minimax errors in a basis which provides a nearly "optimal" sparse decomposition. We briefly summarize the main results. In the following we write $u(d) \sim v(d)$ if there exists $B \geq A > 0$ such that for all $d \in \mathbb{N}$

$$A u(d) \leq v(d) \leq B u(d).$$

A fundamental theorem of Donoho and Johnstone [14] proves that a thresholding estimator (5) for $\lambda = \sqrt{2 \log d}$ satisfies

$$\mathbb{E}(\|P_{\mathbf{V}}f - \hat{f}(g)\|^2) \leq (2 \log d + 1) \left(\sigma^2 + \sum_{n=1}^d \min(|\langle f, \psi_n \rangle|^2, \sigma^2) \right).$$

The upper bound is controlled by the sparsity of the decomposition coefficients $|\langle f, \psi_n \rangle|$. This sparsity is measured by sorting them in decreasing amplitude order. We denote $|\langle f, \psi_{n_r} \rangle|$ the sorted coefficient of rank r : $|\langle f, \psi_{n_r} \rangle| \geq |\langle f, \psi_{n_{r+1}} \rangle|$. If $|\langle f, \psi_{n_r} \rangle|^2 \sim r^{-(\alpha+1)}$ one can verify that

$$\sum_{n=1}^d \min(|\langle f, \psi_n \rangle|^2, \sigma^2) \sim \sigma^{2\alpha/(\alpha+1)}.$$

Suppose that Λ is a ‘‘tail compact’’ set, which means that there exists $C > 0$ and $\epsilon > 0$ such that for all $f \in \Lambda$ and all $M > 0$, we have $\sum_{n \geq M} |\langle f, \psi_n \rangle|^2 \leq C M^{-\epsilon}$. By choosing $d = \sigma^{-\beta}$ for β large enough we then verify with (3) that there exists $C > 0$ such that the maximum error on Λ satisfies

$$\epsilon_M(\sigma) \leq C \sigma^{2\alpha/(\alpha+1)} |\log \sigma|.$$

The last step proves that this upper bound gives the minimax error over Λ up to the $\log \sigma$ term. Donoho [11] proves that if there exists two balls Γ_1 and Γ_2 of a Banach space such that $\Gamma_1 \subset \Lambda \subset \Gamma_2$ and if the sparsity basis $\{\psi_n\}_{n \in \mathbb{N}}$ defines an unconditional basis of this Banach space then there exists $C' > 0$ such that the minimax error satisfies

$$\epsilon_m(\sigma) \geq C' \sigma^{2\alpha/(\alpha+1)}$$

and hence

$$\epsilon_m(\sigma) \leq \epsilon_M(\sigma) \leq C/C' \epsilon_m(\sigma) |\log \sigma|. \quad (6)$$

Neglecting the $\log \sigma$ term, we shall then say that the thresholding estimator is asymptotically minimax. Next section applies these results to wavelet bases which define unconditional bases of all Besov spaces [42].

2.2 Denoising in wavelet orthonormal bases

Wavelet orthonormal bases A wavelet orthonormal basis of $\mathbf{L}^2[0, 1]^2$ is constructed with three mother wavelets ψ^k for $1 \leq k \leq 3$, which are dilated $\psi_j^k(u) = 2^{-j} \psi^k(2^{-j}u)$ at different scales $2^j < 1$, and translated, modulo boundary conditions [39]. It defines a wavelet orthonormal basis

$$\{\psi_j^k(2^j n - u)\}_{j < 0, 1 \leq k \leq 3, 2^j n \in [0, 1]^2}.$$

The inner products of f with translated wavelets can also be rewritten as convolutions sampled at intervals 2^j :

$$\langle f(u), \psi_j^k(2^j n - u) \rangle = f * \psi_j^k(2^j n).$$

A finite dimensional approximation $P_{\mathbf{V}}f$ of f at the scale 2^{j_m} is defined by setting to zero all wavelet coefficients at scales $2^j \leq 2^{j_m}$. It projects f in the space $\mathbf{V}$ of dimension $d = 2^{-2j_m}$ generated by wavelets having a scale $2^j \geq 2^{j_m}$. This space captures the low frequencies of f . It also admits an orthonormal basis of 2^{-2j_m} translated scaling functions $\{\phi_{j_m}(2^{j_m} n - u)\}_{2^{j_m} n \in [0, 1]^2}$ where $\phi_{j_m}(u) = 2^{-j_m} \phi(2^{-j_m}u)$. The Fourier transform of ϕ_{j_m} is mostly concentrated in the low-frequency square $[-2^{-j_m}\pi, 2^{-j_m}\pi]$. The orthogonal projection $P_{\mathbf{V}}f$ is thus also represented by a discrete image of $d = 2^{-2j_m}$ coefficients $f * \phi_{j_m}(2^{j_m} n)$ which are local averages of f in the neighbourhood of each grid point $2^{j_m} n$. This discrete image is typically provided as input measures.

One may also limit the maximum wavelet scale to $2^{j_M} \leq 1$. The 2^{-2j_M} wavelets at larger scales generate the same space as the family of scaling functions $\{\phi_{j_M}(2^{j_M} n - u)\}_{2^{j_M} n \in [0, 1]^2}$. The wavelet representation up to the scale 2^{j_M} of $P_{\mathbf{V}}f$ is thus defined by

$$Wf = \left\{ f * \psi_j^k(2^j n) \right\}_{j_m \leq j \leq j_M, 2^{-j} n \in [0, 1]^2, 1 \leq k \leq 3} \cup \left\{ f * \phi_{j_M}(2^{j_M} n) \right\}_{2^{-j_M} n \in [0, 1]^2}.$$

These discrete wavelet coefficients are computed at scales $2^j \geq 2^{j_m}$ with a fast discrete wavelet transform, with $O(d)$ operations [40]. The amplitudes of wavelet coefficients depend upon the local regularity of f . The more regular f the smaller $|f * \psi_j^k(2^j n)|$ at fine scales 2^j , as quantified by equation (7) for Lipschitz α functions. Figures 1(a) and 1(b) show respectively a piecewise regular image along with its discrete orthogonal wavelet coefficients. Large coefficients are located near edges, where the image is discontinuous. A wavelet thresholding estimator sets to zero all coefficients below a threshold T proportional to the noise standard deviation σ . Setting to zero a wavelet coefficient eliminates a high frequency component, which is equivalent to eliminate the noise with a local averaging operator. This averaging operator adapts to the amplitude of wavelet coefficients and hence to the local signal regularity.




Figure 1: (a) A C^2 -geometrically regular image discretised in dimension $d = 128^2$, (b) 3 levels discrete orthogonal wavelet decomposition of this image. Large wavelet coefficients could clearly be seen along edges. Symlets with four vanishing moments are used [8].

Wavelet thresholding optimality for uniformly Lipschitz models The Lipschitz regularity of a function f in the two-dimensional neighbourhood of a point $u \in \Omega$ is defined by the error of a polynomial approximation. A function $f \in \mathbf{L}^2[0, 1]^2$ is uniformly Lipschitz α on Ω with Lipschitz constant $L > 0$ if for any $u \in \Omega$ there exists a polynomial q_u of degree smaller than α such that

$$\forall u' \in \Omega, |f(u') - q_u(u')| \leq L |u - u'|^\alpha.$$

The minimal value of L defines the norm $\|f\|_{C^\alpha}$ of f in the Hölder space $\mathbf{C}^\alpha$ over Ω . The set $\Lambda \subset \mathbf{L}^2(\Omega)$ of all functions f which are Lipschitz α for a constant L defines a determinist model of images corresponding to a ball in this Hölder space.

We shall suppose that wavelet ψ^k has m vanishing moments, which means that for any polynomial $q(u)$ of degree m we have

$$\int q(u) \psi^k(u) du = 0.$$

Wavelet coefficients provide a characterization of Hölder spaces $\mathbf{C}^\alpha[0, 1]^2$ of uniformly Lipschitz α functions. For $\alpha \leq m$, one can prove [26, 42] that $f \in \mathbf{C}^\alpha[0, 1]^2$ if and only if there exists $C > 0$ such that for all $2^j n \leq [0, 1]^2$

$$|f * \psi^k(2^j n)| \leq C 2^{(\alpha+1)j}. \quad (7)$$

One can verify that this inequality implies that sorted wavelet coefficients of f of rank r decay at most like $r^{-(\alpha+1)/2}$. Wavelets define unconditional bases of Hölder spaces [42]. It results from (6) that a wavelet thresholding estimator has a maximum error $\epsilon_M(\sigma)$ which is asymptotically minimax with

$$\epsilon_m(\sigma) \sim \sigma^{2\alpha/(\alpha+1)}.$$

In that sense a wavelet thresholding estimator is nearly optimal to suppress noise for all levels of uniform regularity $\alpha \leq m$ and it adapts to the regularity exponent α which may be unknown. Such results can be extended to ball of any Besov space [15] which gives other conditions to specify the regularity of f .

Wavelet sub-optimality over $\mathbf{C}^\alpha$ geometric models Along an edge curve, image values have regular variations when moving parallel to the edge curve but they may be discontinuous when crossing the edge boundary. The image regularity is therefore anisotropic. We consider models of $\mathbf{C}^\alpha$ geometrically regular images, whose definition was introduced in Peyré and Mallat [45], inspired by Donoho [12], Korostelev and Tsybakov [30].

Definition 2.1. (geometrically $\mathbf{C}^\alpha$) A function $f \in \mathbf{L}^2[0, 1]^2$ is said to be geometrically $\mathbf{C}^\alpha$ with Lipschitz constant L if f is uniformly Lipschitz α with constant L over $[0, 1]^2 - \{\gamma_i\}_{1 \leq i \leq n}$, where the γ_i are uniformly Lipschitz α curves with constant L , and these curves do not intersect tangentially.

The set $\Lambda \subset \mathbf{L}^2[0, 1]^2$ of all functions f which are geometrically $\mathbf{C}^\alpha$ with constant L defines a model of images which incorporate edges along regular curves, which may intersect at junctions. Wavelet coefficients have a large amplitude along edges, as illustrated in Figure 1(b). A wavelet thresholding algorithm keeps the noise added to these large amplitude wavelet coefficients, which dominate the overall error of a wavelet thresholding estimator. One can verify that the resulting maximum error over a set of $\mathbf{C}^\alpha$ geometrically functions satisfies

$$\sup_{f \in \Lambda} \mathbb{E}(\|\hat{f} - f\|^2) \sim \sigma |\log \sigma|. \quad (8)$$




Figure 2: (a): the noisy image is the sum of a $\mathbf{C}^2$ geometrically regular image ranging in $[-1, 1]$ with an additive Gaussian white noise of variance $\sigma^2 = 0.005$. (b): orthogonal wavelet coefficients of the noisy image computed on 3 scales 2^j . Larger coefficients are darker. (c): non-zero wavelet coefficients after soft thresholding are shown in black. (d): denoised estimated image recovered from thresholded wavelet coefficients. The PSNR is 28.97dB for the noisy image and 33.29dB for its denoised counterpart.

It does not depend upon the Lipschitz exponent α and only takes advantage of the fact that edges have a finite length, which limits the number of large wavelet coefficients. A similar error is obtained over a much larger set of bounded variation functions [39]. Figures 2(a) to 2(d) gives an example of denoising by thresholding the wavelet orthogonal coefficients of a geometrically regular image. As expected, the thresholding retains wavelet coefficients in the neighbourhood of edges and hence retains noise in this neighbourhood. The soft thresholding attenuates the amplitude of the noise at the cost of restoring edges which are smoothed out. Figure 3(a) shows the mean-squared error $\log \epsilon_{ms}(\sigma)$ as a function of $\log \sigma$ for these $\mathbf{C}^\alpha$ geometrically regular images. The slope does not depend upon α and remains equal to 1. The implementation was done with a nearly translation invariant thresholding estimator obtained by averaging translated orthogonal wavelet estimators [7] (see D.2 for details). This averaging reduces the error by a multiplicative factor but does not modify the asymptotic slope. Numerical errors are computed on average over a random process whose realizations are $\mathbf{C}^\alpha$ geometrically regular functions, whose edge length concentrates as explained in D.1. Estimation errors thus remain of the same order of magnitude as the maximum error for all typical realizations.

2.3 Optimal denoising of geometrically regular images

To reduce the error of wavelet thresholding estimators we must take advantage of edges regularity, these edges being $\mathbf{C}^\alpha$ curves. Korostelev and Tsybakov [30] proved that the minimax rate over $\mathbf{C}^\alpha$ geometrically regular images remains

$$\epsilon_m \sim \sigma^{2\alpha/(\alpha+1)}, \quad (9)$$

which is identical to the minimax rate for functions which are uniformly Lipschitz α without edges. In other words, the presence of edges does not affect the minimax denoising rate if these edges are themselves $\mathbf{C}^\alpha$ regular curves which intersect over a finite number of junctions. However, this result was not constructive. An important research effort has been devoted to reach this minimax rate through sparse decompositions in appropriate bases or frames [4, 45, 10, 32]. However, much better results have been obtained with deep convolutional neural networks. We briefly review these results to understand where are the difficulties.




Figure 3: (a): Mean-squared error $\epsilon_{ms}(\sigma)$ of a wavelet thresholding estimator, as a function of the noise variance σ^2 , computed on average over C^α geometrically regular images defined in D.1. Red to blue curves correspond to Lipschitz exponent $\alpha \in \{2, 1.5, 1.2, 1\}$. The slope of $\log \epsilon_{ms}$ remains equal to 1 as a function of $\log \sigma$. (b): same error calculations for a convolutional neural network estimator computed with a UNet. For each Lipschitz exponent α , $\log \epsilon_{ms}$ nearly superimposes with a straight line of slope $2\alpha/(\alpha + 1)$ which corresponds to the minimax rate.




Figure 4: (a) An edge creates large wavelet coefficients whenever the square wavelet supports intersect the edge and the wavelets oscillate across the edge. The support width is proportional to the scale 2^j so the total number of large wavelet coefficients along an edge is proportional to 2^{-j} . (b) Curvelets, shearlets and bandlets have elongated supports whose width is proportional to 2^j but whose length is proportional to $2^{j/\alpha}$, with $\alpha = 2$ for curvelets and shearlets. If the edge is a C^α curve then the number of non-negligible coefficients is proportional to $2^{-j/\alpha}$.

Edges of C^α geometrically regular images have two important properties. They correspond to pointwise discontinuities and these discontinuities propagate along regular C^α curves. To avoid retaining noise in the neighbourhood of edges, an edge is covered with few elongated vectors with a support which is as narrow as possible. The length of these vectors depends upon the edge regularity as illustrated in Figure 4(b). This strategy reduces the number of vectors needed to cover the edge compared to the number of wavelets illustrated in Figure 4(a), but it requires to adapt the support of the frame vector to the edge regularities. Curvelets, shearlets and bandlets adopt this strategy.

Curvelet frames Curvelet frames [4, 6] have been constructed to obtain optimal denoising of images whose edges are closed C^2 contours, without corners. Curvelets have an elongated rectangular support of width 2^j and length $2^{j/2}$. A curvelet frame is constructed by rotating this support and translating it on a rotated anisotropic grid whose intervals is proportional to 2^j and $2^{j/2}$. Whereas C^2 edge curves create of the order of 2^{-j} large wavelet coefficients at each scale, they only produce of the order of $2^{-j/2}$ large curvelet coefficients. A curvelet representation of such edge curves is therefore much sparser than a wavelet basis representation. A curvelet denoising estimator is computed by thresholding the curvelet coefficients of a noisy image [5]. The resulting estimation error over C^2 edge curves reaches the minimax bound (9) for $\alpha = 2$ if there is no corner. This important theoretical result has limited applications because curvelet become suboptimal when $\alpha \neq 2$ and when the edges join at corners. Shearlets [23, 32] provide efficient implementations of curvelets for discretised images, by replacing rotation operators by shearing operators along the horizontal and vertical directions.

Adaptivity to geometric regularity with bandlets To build optimally sparse representations of images with edges, the support of basis vectors must be adapted to the regularity of edge curves. If the edges are C^α curves then, at a scale




Figure 5: This figure shows the denoising performance of a UNet trained on two datasets. For the top row image, the UNet was trained on the LSUN bedrooms dataset [57] in dimension $d = 128^2$ pixels. On the second row it was trained on C^α geometrically regular images of same size. The image amplitudes are normalized in $[-1, 1]$. The first column shows an original test image which does not belong to the training set. The second column shows a noisy image with a noise variance $\sigma^2 = 1$, which corresponds to a PSNR of 6dB. The third column displays the denoised image computed by the UNet. The PSNR is 22.03 dB for the bedroom image 30.03dB for the C^α geometrically regular image. The fourth column shows a zoom near edges. In both case the estimator reproduces very well the image edges and corners.

2^j , a basis vector of width 2^j must have a length of the order of $2^{j/\alpha}$, as shown in Figure 4(b). It means that different basis vectors must be used when the geometric regularity changes. This problem was solved by bandlet dictionaries, which include an infinite number of orthonormal bases. These bases are obtained by transforming orthogonal wavelet coefficients with adapted unitary operators [45]. A best basis, regrouping vectors having the appropriate geometry is selected with a best basis selection algorithm [34, 45]. A bandlet denoising estimator selects such a basis from the noisy signal and thresholds the image coefficients in the selected best basis [18]. It has a maximum error over C^α geometrically regular images which reaches the minimax bound (9) for all α , despite the presence of corners. The difficulty of this approach is that the size of the dictionary grows quickly with the diversity of geometric image properties.

Denoising with deep neural networks Despite interesting asymptotic results, the effort to denoise image with geometrically adaptive bases has not been so effective on natural images, which have a complex geometry. Non-linear estimations with deep neural networks provide much better results and seem to follow very different principles. These neural networks involve cascades of filters with a fixed short support, typically of size 3×3 , and pointwise non-linearities such as *ReLU*. They do not have filters with elongated supports, adapted to the image geometry. Still, they produce much smaller denoising error. Deep neural networks are trained to compute denoising estimators $\hat{f}(g)$ by minimising a mean-squared error $\mathbb{E}_{P_{\mathbf{V}f,g}}(\|\hat{f}(g) - P_{\mathbf{V}f}\|^2)$ where the expected value is computed over the noise distribution and over a set of training data. Impressive denoising results are obtained over images which do not belong to the training set, which considerably improve previous state of the art results [19, 53]. Figure 5 shows denoising examples trained over a dataset of natural images (the LSUN bedroom dataset [57]) and over a dataset of piecewise regular images. The deep neural network has a UNet architecture [49]. In both cases, these estimators retrieve the image geometry, even at very low signal to noise ratios. The network estimation restores sharp contours as soon as long enough portions of these contours emerge from the noise. Objects delimited by edges which are not long enough or not sufficiently contrasted (given the noise level) are not restored. This is the case for the painting in the top right corner of the LSUN bedroom example. Figure 3(b) shows the mean-squared denoising error $\epsilon_{ms}(\sigma)$ of this UNet over a set of C^α geometrically regular images defined in D.1. This mean-squared error reaches the optimal minimax rate as a function of σ , for all $1 \leq \alpha \leq 2$. This had already been observed in [28]. However, there is currently no mathematical understanding of the performance of deep neural network estimators.

3 Denoising with a scattering transform

The previous section shows that curvelets, shearlets and bandlets can define optimal estimators of C^α geometrically regular images by taking advantage of their directional regularity to produce a sparse representation. Numerical results also show that deep neural networks can define better estimators having asymptotically optimal denoising performances over C^α geometrically regular images. As opposed to curvelets, Xlets or bandlets, they do not involve linear filters having elongated support in multiple directions but they are computed with a cascade of linear filters and pointwise non-linearities. A scattering transform has been introduced in [41, 2] as a simplified model of a two layers convolutional neural network, with complex wavelet filters and a modulus non-linearity. Section 3.3 shows numerically that an estimator regularized by scattering ℓ^1 norms can achieve an asymptotically minimax denoising over C^α geometrically regular images, which adapts to α . We state a mathematical conjecture which formalises these numerical observations, opening a mathematical bridge with deep neural network estimators. Although we do not prove the conjecture, we give intermediate mathematical results which partly explains these observations. Section 3.1 begins by a brief review of complex directional wavelet transforms. Section 3.2 proves that the image regularity is partly captured by a sparse set of scattering coefficients. However, scattering coefficients do not involve elongated filters such as curvelets or bandlets, whose support are adapted to the edge regularity. We show that this is replaced by another subset of scattering coefficients, which capture the edge singularity with large ℓ^1 norms.

3.1 Directional complex wavelet transform

Wavelets can capture anisotropic regularities if they have directional vanishing moments aligned with the direction where the image is regular. Such wavelets can be created by rotating a single mother wavelet [39]. In the following we review the properties of a dyadic directional wavelet transform and prove that the regularity of a contour can be captured by a subset of wavelets coefficients. We write $\|f\|_1 = \int |f(u)| du$ the L^1 norm of a function. For simplicity, we shall sometime loosely address these L^1 norms as ℓ^1 norms because they are numerically computed as ℓ^1 norms of functions discretised on a sampling grid.

Dyadic directional wavelet transform We define 4 wavelets $\psi^k(u) = \psi(r_k u)$ for $u \in \mathbb{R}^2$ which are rotations of a single complex wavelet ψ , where r_k is a rotation in $\mathbb{R}^2$ by an angle $-k\pi/4$ for $0 \leq k < 4$. A dyadic wavelet transform of f is computed as convolutions with dilated wavelets $\psi_j^k(u) = 2^{-2j} \psi^k(2^{-j} u)$ for $2^j \leq 2^{j_M}$, without subsampling. The multiplicative factor 2^{-2j} normalises the L^1 norm of all wavelets

$$\|\psi_j^k\|_1 = \int |\psi_j^k(u)| du = \int |\psi(u)| du = 1.$$

At the largest scale 2^{j_M} , the lower frequencies are retained by a convolution with a scaling function $\phi_{j_M}(u) = 2^{-2j_M} \phi(2^{-j_M} u)$, where $\phi \geq 0$ is positive and $\int \phi(u) du = 1$:

$$Wf = \left(f * \phi_{j_M}, f * \psi_j^k \right)_{0 \leq k < 4, j \leq j_M}. \quad (10)$$

Such a dyadic wavelet transform is invertible and has a stable inverse [39] if the wavelet satisfies Littlewood-Paley inequalities:

$$\exists \delta < 1, \forall \omega \neq 0, 1 - \delta \leq |\hat{\phi}_{j_M}(\omega)|^2 + \sum_{j=-\infty}^{j_M} \sum_{k=0}^3 |\hat{\psi}_j^k(\omega)|^2 \leq 1 + \delta. \quad (11)$$

We define wavelet transforms in $[0, 1]^2$ with a standard periodization of wavelets in $\mathbb{R}^2$ [40]. In numerical computations, the image is known at a finite resolution 2^{j_m} over a discrete grid with an orthogonal projection on a space $\mathbf{V}$. Section 2.2 explains that it is equivalent to limit all wavelet scales to $2^j \geq 2^{j_m}$. All convolutions are performed on this sampling grid, without subsampling. Figure 6(b) shows the real and imaginary parts of $f * \psi_j^k$ along 4 orientations and 3 dyadic scales, for f being a geometrically regular function shown in Figure 6(a).

To capture directional image regularities, we shall impose that ψ has m directional vanishing moments in directions $\theta \in [-\pi/4, \pi/4]$. It means that wavelets rotated in these directions are orthogonal to any one-dimensional polynomial function $u_1 \mapsto q(u_1)$ of degree $m - 1$:

$$\forall \theta \in [-\pi/4, \pi/4], \forall u_2 \in [0, 1], \int q(u_1) \psi(r_\theta(u_1, u_2)) du_1 = 0.$$

A gives a sufficient condition on the Fourier transform of wavelets to obtain directional vanishing moments. Figure 7 shows the real and imaginary parts of a wavelet having $m = 2$ vanishing moments in directions $\theta \in [-\pi/4, \pi/4]$, along




Figure 6: Dyadic directional wavelet transform of a C^2 geometrically regular image, discretised in dimension $d = 128^2$. (a): Original image. (b): Real and imaginary parts of wavelet coefficients $f * \psi_j^k$ of image (a), for scales 2^j and orientations k which respectively vary vertically and horizontally. (c): Modulus $|f * \psi_j^k|$ of these wavelet coefficients. They are large when the vanishing moments of ψ_j^k are not aligned with the tangent to the edge. Plots (b) and (c) are normalized such that images on a same row share the same minimum and maximum value.




Figure 7: Real (a) and imaginary parts (b) of a wavelet ψ having two vanishing moments in directions $\theta \in [-\pi/4, \pi/4]$. Figures (c) and (d) show respectively ψ 's modulus and phase (where $|\psi|$ is non negligible). One should notice that lines of constant phase are nearly vertical.

with its modulus and phase. This latter oscillates in the vanishing moments directions and remains nearly constant in the orthogonal direction.

If f is uniformly Lipschitz $\alpha \leq m$, since each wavelet ψ_j^k has m vanishing moments, one can verify as in (7) (modulo the difference of normalisation) that there exists $C > 0$ such that

$$\forall u \in [0, 1]^2, |f * \psi_j^k(u)| \leq C 2^{j\alpha}. \quad (12)$$

For C^α geometrically regular images, Section 2.2 explains that wavelet coefficients are large along the edge where f is discontinuous. Standard wavelets thus fail to exploit the regularity of contours, yielding representations that are less sparse than theoretically possible and consequently suboptimal denoising performance. In what follows, we demonstrate that this geometric regularity can be effectively captured using wavelets with vanishing moments in the direction tangent to the edge. Since the image intensity varies regularly in this direction, the corresponding wavelet coefficient is small. This is formalized by the following theorem, which considers a periodic image having a single edge whose tangents have an angle in $[(k-1)\pi/4, (k+1)\pi/4]$.

Theorem 3.1. *Suppose that ψ is a C^2 function with two vanishing moments and two directional vanishing moments in the cone of angles $[-\pi/4, \pi/4]$, such that $|\psi|$ decays faster than any rational function. Let $\alpha \in [1, 2]$. Let f be a periodic image on $[0, 1]^2$ which is uniformly Lipschitz α above and below an edge curve, which is itself Lipschitz α and whose tangents have an angle in $[(k-1)\pi/4, (k+1)\pi/4]$. For all $s < \alpha$, there exists C such that:*

$$\forall j \leq j_M, \forall u \in [0, 1]^2, |f * \psi_j^k(u)| \leq C 2^{(s-1)j}. \quad (13)$$

This theorem proves that a wavelet having vanishing moments in the direction of an edge tangent can take advantage of the edge regularity to produce small amplitude coefficients. On the contrary, wavelets having no directional vanishing moments in these directions produce large amplitude wavelet coefficients. It appears in Figure 6(c) which shows that $|f * \psi_j^k|$ has a large amplitude along edges where the tangents have an angle sufficiently different from $k\pi/4$.

3.2 Scattering transform for geometric regularity

To define minimax denoising estimators for C^α geometrically regular functions, we need to precisely capture the regularity of these functions. It has three components. Outside edges, the functions are uniformly Lipschitz α , which is captured by the fast decay of wavelet coefficients $|f * \psi_j^k|$. Theorem 3.1 proves that anisotropic regularity of f along edges is partly captured by the fast decay of $|f * \psi_j^k|$ when the wavelet directional vanishing moments are aligned with the edge tangent. However, the other wavelet coefficients are unable to characterize the sharp profile of the edge, which is a pointwise discontinuity when we move in the direction orthogonal to the edge curve. This issue is addressed by a scattering transform, which applies a second wavelet transform on $|f * \psi_j^k|$ to simultaneously capture the edge regularity and its sharp profile. A wavelet scattering transform applies a second wavelet transform to each $|f * \psi_j^k|$

$$Sf = \left(f * \phi_{j_M}, W|f * \psi_{j,k}| \right)_{j \leq j_M, k < 4}$$

with

$$W|f * \psi_j^k| = \left(|f * \psi_j^k| * \phi_{j_M}, |f * \psi_j^k| * \psi_{j'}^{k'} \right)_{j' \leq j_M, k' < 4}.$$

Real and imaginary parts of such coefficients are shown in Figure 8(b) for a fixed $|f * \psi_j^k|$, shown in Figure 8(a), along 4 different orientations k' and 3 dyadic scales $j' > j$. We compute the L^1 norm of each of these second order wavelet coefficients. Since $\phi_{j_M} \geq 0$ and $\int \phi_{j_M}(u) du = 1$, it results that $\| |f * \psi_j^k| * \phi_{j_M} \|_1 = \|f * \psi_j^k\|_1$. The scattering coefficients L^1 norm can thus be written

$$\|W|f * \psi_j^k|\|_1 = \left(\|f * \psi_j^k\|_1, \| |f * \psi_j^k| * \psi_{j'}^{k'} \|_1 \right)_{j' \leq j_M, k' < 4}.$$

Each $|f * \psi_j^k| * \psi_{j'}^{k'}$ computes the variations of $|f * \psi_j^k|$ in a direction $\pi k'/4$, in neighbourhoods of size proportional to $2^{j'}$. If ψ is regular with m derivatives then Mallat [41] prove that $\| |f * \psi_j^k| * \psi_{j'}^{k'} \|_1$ becomes negligible when $2^{j'} < 2^j$ because $|f * \psi_j^k|$ has a Fourier transform concentrated at lower frequencies than $\psi_{j'}^{k'}$. We thus only compute scattering coefficients for $j' > j$ as Bruna and Mallat [2]. These scattering coefficients are shown in Figure 8(c).

Directional regularity appears through small L^1 norms For C^α geometrically regular images, the amplitude of $|f * \psi_j^k|$ has a fast decay at small scales 2^j besides the neighbourhoods of edges whose tangents are not aligned with ψ_j^k vanishing moments, i.e. when the edge tangent has an angle which is not in $[(k-1)\pi/4, (k+1)\pi/4]$. Applying




Figure 8: (a): Wavelet transform modulus $|f * \psi_j^k|$ for $j = j_m$ and $k = 0$. (b): Scattering coefficients $|f * \psi_j^k| * \psi_{j'}^{k'}$ for $j' > j$ varying over three dyadic scales and 4 orientations indexed by k' . (c): Modulus $||f * \psi_j^k| * \psi_{j'}^{k'}|$ of the previous scattering coefficients. Large amplitudes are obtained near edges which are oriented neither along $k\pi/4$ nor $k'\pi/4$. Coefficient maps from the third column, corresponding to $k' = k^\perp$, exhibit large amplitudes only in the vicinity of corners. Plots (b) and (c) are normalized so that images on a same row share the same minimum and maximum value.

ψ_j^k therefore only partially sparsifies f , since on average half of the large wavelet coefficients remain (as contours are equally distributed across all orientations). In this paragraph, we show how to overcome this alignment issue by applying a wavelet $\psi_{j'}^{k^\perp}$, with $k^\perp = (k+2) \bmod 4$, to $|f * \psi_j^k|$. Indeed, large amplitude values of $|f * \psi_j^k|$ are located at edge points where the direction of the tangent is in the complementary interval $[(k+1)\pi/4, (k+3)\pi/4]$ modulo π . They vary regularly along the edge curve whose tangent is within this angular interval. However, this interval corresponds to the directions where $\psi_{j'}^{k^\perp}$ has vanishing moments, as illustrated in Figure 9. With the same argument as in Theorem 3.1, a convolution with $\psi_{j'}^{k^\perp}$ will produce small amplitude coefficients everywhere outside junctions between edges, which define corners. The following theorem derives an upper bound on $||f * \psi_j^k| * \psi_{j'}^{k^\perp}||_1$ which decreases to zero when the second scale $2^{j'}$ decreases, for all $1 \leq \alpha \leq 2$. This result is obtained by proving that large coefficients at junctions do not affect the asymptotic decay of the L^1 norm.

Theorem 3.2. *Suppose that ψ satisfies the same hypothesis as in Theorem 3.1. If f is a C^α geometrically regular function with $1 \leq \alpha \leq 2$ then, for all $s < \alpha$, there exists C such that:*

$$\forall j < j' \leq j_M, \forall k < 4, ||f * \psi_j^k| * \psi_{j'}^{k^\perp}||_1 \leq C 2^{sj'}, \quad (14)$$

and

$$\forall s < \alpha, \forall k < 4, \sum_{j=-\infty}^{j_M} \sum_{j'=j+1}^{[j/\alpha]} 2^{-sj'} ||f * \psi_j^k| * \psi_{j'}^{k^\perp}||_1 < +\infty. \quad (15)$$




Figure 9: Wavelet coefficients $|f * \psi_j^k|$ are large at edge locations where ψ_j^k , shown as small squares, oscillate across the edge discontinuity. These coefficients have the same geometric regularity as the edge. Their wavelet transform $|f * \psi_j^k| * \psi_{j'}^{k\perp}$ have a small amplitude because the larger scale wavelet $\psi_{j'}^{k\perp}$, shown as a large square, oscillates along the perpendicular direction, i.e. along the edge tangent, where the edge is regular.

The proof is in C. It shows that $\|f * \psi_j^k * \psi_{j'}^{k\perp}\|$ is small where f is uniformly Lipschitz α because $|f * \psi_j^k|$ is small. This quantity remains small in the neighbourhood of edge points which are not corners, because either ψ_j^k or $\psi_{j'}^{k\perp}$ can take advantage of the directional image regularity along the edge tangent. These coefficients may however become large at corners where the image has no directional regularity, but since there is a finite number of such corners and since wavelets have a square support, they do not affect the asymptotic decay of the L^1 norm. This is illustrated by Figure 8(c) which shows that the only large amplitudes of $\|f * \psi_j^k| * \psi_{j'}^{k\perp}\|$ correspond to edge junctions. The result and the proof remain valid if the modulus $|f * \psi_j^k|$ is replaced by a rectification $ReLU(f * \psi_j^k)$ of the real and imaginary parts, where $ReLU(a) = \max(a, 0)$ is the rectifier used in most convolutional networks.

As explained in Section 2.1, the decay rate of decomposition coefficients in a frame is directly related to the denoising error obtained with a sparse prior (4). Theorem 3.2 specifies the decay of second order wavelet coefficient amplitudes, which depends on α despite the presence of discontinuities along contours. Nevertheless, this decay is not sufficient to get an optimal denoiser by replacing the sparse prior of equation (4) with a combination of scattering L^1 norms $(\|f * \psi_j^k| * \psi_{j'}^{k\perp}\|_1)_{j \leq j' \leq j_M, k < 4}$. Indeed, we are missing the other scattering coefficients $(|f * \psi_j^k| * \psi_{j'}^{k'})_{j \leq j' \leq j_M, k' \neq k^\perp}$ to get an invertible representation of $(|f * \psi_j^k|)_{j \leq j_M, k < 4}$. For a C^α geometrically regular function, these missing coefficients are not sparse but are concentrated in the neighbourhood of edge curves. To obtain an optimal denoising estimator, we must enforce this property with appropriate conditions on these scattering coefficients.

Edge profile captured by large L^1 norms The edges of C^α geometrically regular functions correspond to sharp pointwise discontinuities if we move in the direction perpendicular to their tangents. Section 2.3 explains that it requires to adapt the support of curvelets, shearlets and bandlets to obtain a sparse representation. Wavelet coefficients $|f * \psi_j^k|$ have a large amplitude along the edge, with a narrow support when moving away from the edge, which is equal to the support of ψ_j^k and hence proportional to 2^j . We shall see that this property is captured by imposing that scattering coefficients $|f * \psi_j^k| * \psi_{j'}^{k'}$ have a sufficiently large L^1 norm for $k' \neq k^\perp$. It appears through Young's inequality, which implies that

$$\| |f * \psi_j^k| * \psi_{j'}^{k'} \|_1 \leq \|f * \psi_j^k\|_1 \| \psi_{j'}^{k'} \|_1. \quad (16)$$

Since $\| \psi_{j'}^{k'} \|_1 = 1$ and $|f * \psi_j^k| \geq 0$,

$$\|f * \psi_j^k\|_1 \| \psi_{j'}^{k'} \|_1 - \| |f * \psi_j^k| * \psi_{j'}^{k'} \|_1 = \int \left(\int |f * \psi_j^k(u)| | \psi_{j'}^{k'}(v - u) | du - \left| \int |f * \psi_j^k(u)| \psi_{j'}^{k'}(v - u) du \right| \right) dv \geq 0. \quad (17)$$

This difference is minimised if, for all v , $|f * \psi_j^k|(u)$ is small except for values of u such that $\psi_{j'}^{k'}(v - u)$ has a nearly constant phase or is negligible. This property is studied by Bruna and Mallat [3] to characterize point processes. For directional wavelets shown in Figure 7, $|\psi|$ is nearly isotropic and the lines of constant phase are nearly vertical. Since k' is a rotation index, the domain where $|\psi_{j'}^{k'}|$ is non-negligible does not depend upon k' . Over this domain the lines of constant phase are nearly vertical straight lines rotated by $k'\pi/4$ and thus have an angle $k'\pi/4 + \pi/2$. Remember that $|f * \psi_j^k|$ has a large amplitude near edge points where the edge tangent has an angle close to $k\pi/4 + \pi/2$. If $k' = k$, lines of constant phase have thus an angular discrepancy of, at most, $\pi/8$ with edges' tangents.

Given a fixed $\|f * \psi_j^k\|_1$, maximising $\| |f * \psi_j^k| * \psi_{j'}^{k'} \|_1$ requires to concentrate the support of $|f * \psi_j^k|$ on regions where the phase of $\psi_{j'}^{k'}$ varies as little as possible. Section 3.3 shows empirically that such regions could adapt to the tangent, despite the potential misalignment between lines of constant phase and edge tangents. Indeed, as explained in Section

3.3, $\|f * \psi_j^k * \psi_{j'}^{k'}\|_1$ is jointly optimized with a data fidelity term which forces f to be close from the noisy observation. In particular, for j' big enough compared to the noise variance σ , tangents are constrained by this observation.

For a fixed edge tangent close to $k\pi/4 + \pi/2$, $\|f * \psi_j^k * \psi_{j'}^{k'}\|_1$ provides a measure of $|f * \psi_j^k|$'s support width. If $k' = k \pm 1$, $\|f * \psi_j^k * \psi_{j'}^{k'}\|_1$ is small when the edge tangent has an angle close to $k'\pi/4$, thanks to $\psi_{j'}^{k'}$ vanishing moments (the proof is nearly identical to that of Theorem 3.2). If not, $\|f * \psi_j^k * \psi_{j'}^{k'}\|_1$ behaves as $\|f * \psi_j^k * \psi_{j'}^k\|_1$, providing big coefficients whose amplitudes measure $|f * \psi_j^k|$'s support width. Such coefficients dominate in $\|f * \psi_j^k * \psi_{j'}^{k'}\|_1$. Maximising these $\mathbf{L}^1$ norm thus still concentrates the support width of $|f * \psi_j^k|$, but not as strongly as $\|f * \psi_j^k * \psi_{j'}^k\|_1$ since the lines of constant phase are further away of the edge tangent.

These properties show that the amplitude of $\|f * \psi_j^k * \psi_{j'}^{k'}\|_1$, for $k' \neq k^\perp$, provides a measure of the concentration of wavelet coefficients $|f * \psi_j^k|$ around the edge curve, by taking advantage of cancellation effects due to the phase of $\psi_{j'}^{k'}$. A parallel can be drawn with the phenomenon of activation maximization observed in neural networks. Whereas the first layer essentially encodes wavelet filters [38], in deeper layers it has been shown that the maximisation of neuron activations corresponds to the detection of important image patterns [58, 47]. Scattering coefficients correspond to neurons in the second hidden layer, computed with wavelet filters. The maximisation of $\|f * \psi_j^k * \psi_{j'}^{k'}(u)\|_1$ for $k' \neq k^\perp$ corresponds to the detection of an edge point at u . An outstanding issue is to relate the values of $\|f * \psi_j^k\|_1 - \|f * \psi_j^k * \psi_{j'}^{k'}\|_1$ to precise mathematical properties of edge profiles.

3.3 Numerics and conjecture on the minimax optimality of scattering denoising

This section shows numerically that the scattering $\mathbf{L}^1$ norms do provide asymptotically minimax estimations of $\mathbf{C}^\alpha$ geometrically regular images contaminated by additive noise and gives a mathematical conjecture. A noisy discretised signal is defined in equation (1) by an orthogonal projection $P_{\mathbf{V}}g = P_{\mathbf{V}}f + Z$ of a function $f \in \Lambda$, where Z is a Gaussian white noise of variance σ^2 . An estimation $\hat{f}(g)$ of f is calculated by finding the solution of a variational problem

$$\hat{f}(g) = \operatorname{argmin}_{h \in \mathbf{V}} \frac{1}{2} \|h - P_{\mathbf{V}}g\|^2 + \sigma^2 U(h).$$

If U is defined by $\mathbf{L}^1$ norms of wavelet coefficients, then Section 2.2 shows that it is suboptimal to estimate $\mathbf{C}^\alpha$ geometrically regular images. This section verifies it numerically with a dyadic wavelet transform. We then show numerically that if this denoising estimator is also regularized by $\mathbf{L}^1$ norms of scattering coefficients, then it reaches the asymptotic minimax rate for all $1 \leq \alpha \leq 2$. We state a mathematical conjecture on the minimax optimality of such scattering denoisers.

Dyadic wavelet transform denoising A dyadic wavelet transform denoising is computed from the $\mathbf{L}^1$ norms of wavelet coefficients, limited to the finest discretisation scale 2^{j_m} , up to a larger scale 2^{j_M} :

$$U(h) = \lambda \sum_{j=j_m}^{j_M} \sum_{k=0}^3 2^{-j} \|h * \psi_j^k\|_1. \quad (18)$$

In numerical computations, $h * \psi_j^k$ is calculated with a discrete convolution over the discretisation grid and the $\mathbf{L}^1$ norm is computed with an ℓ^1 norm over this sampling grid. This estimator can be interpreted as a translation invariant wavelet denoising estimator [7]. The normalisation factor 2^{-j} is different from wavelet orthonormal bases because dyadic wavelets ψ_j^k have an $\mathbf{L}^1$ normalisation and the wavelet transform is not subsampled.

The mean-squared error $\epsilon_{ms}(\sigma) = \mathbb{E}_{f,g}(\|\hat{f}(g) - f\|^2)$ is computed on average over a set Λ of $\mathbf{C}^\alpha$ geometrically regular images, defined in D.1, for $\alpha = 2$. Typical images in this set have a fixed Lipschitz bound and an edge length which concentrates so that the mean-squared error is of the order of the maximum error. Images are discretised over a grid of dimension $d = 128^2$. The $\mathbf{C}^\alpha$ images in Λ are generated following the algorithm proposed by Kadkhodaie et al. [28] to sample random edges which are $\mathbf{C}^\alpha$ with a bounded Lipschitz constant. The algorithm is further described with examples of images in D.1. For each value of α , the mean-squared error ϵ_{ms} is computed on average over 100 realizations in Λ .

We optimize numerically λ in order to minimise ϵ_{ms} , which gives $\lambda = 1.2$. Figure 10(a) shows in red the value of the mean-squared error $\log \epsilon_{ms}(\sigma)$ as a function of $\log \sigma$, which has a slope of 1. Although $\epsilon_{ms} \leq \epsilon_M$ it remains of the same order because it is computed over images having edges of nearly the same length and which therefore yield an estimation error which is close to the maximum error $\epsilon_M(\sigma)$. The slope is compatible with the theoretical bound

computed over orthogonal wavelet estimators (8), where $\epsilon_M(\sigma) \sim \sigma |\log \sigma|$. The effect of $\log \sigma$ term is difficult to detect numerically. However, the numerics shows that the error rate is well above the minimax error $\epsilon_m(\sigma) \sim \sigma^{4/3}$ for $\alpha = 2$.

Scattering denoising A scattering transform denoising is computed from the $\mathbf{L}^1$ norms of scattering coefficients. According to the analysis of Section 3.2, the geometric regularity of f is specified by minimising the $\mathbf{L}^1$ norms $\|f * \psi_j^k\|_1$ which depend on the total length of the edge curves, by minimising $\|f * \psi_j^k\|_1$ which specifies the directional image regularity parallel to the edge curves, and by maximising $\|f * \psi_j^k\|_1$ for $k' \neq k^\perp$ which specifies the edge profile across the edge. This joint optimization is computed by defining a regularity metric which is a weighted sum of these 3 terms, with a negative sign for the $\mathbf{L}^1$ norms that must be maximised

$$U(h) = \sum_{j=j_m}^{j_M} \sum_{k=0}^3 \left(\lambda 2^{-j} \|h * \psi_j^k\|_1 + \gamma \sum_{j'>j}^0 2^{-j'} \|h * \psi_j^k\|_1 - \sum_{k' \neq k^\perp} \eta_{k'-k} \sum_{j'>j}^0 2^{-j'} \|h * \psi_j^k\|_1 \right) \quad (19)$$

As in the dyadic wavelet transform case, experiments are performed over sets of $\mathbf{C}^\alpha$ geometrically regular images with corners, for different $1 \leq \alpha \leq 2$. We optimize numerically the constants $\lambda, \gamma, \eta_0, \eta_{\pm 1}$ in order to minimise the mean-squared error $\log \epsilon_{ms}$ over a range of noise values σ , with an optimization procedure discussed in D.2. We impose that $\eta_1 = \eta_{-1}$ because ψ_j^{k+1} and ψ_j^{k-1} have a symmetric role relatively to $|f * \psi_j^k|$. The computed optimized constants are independent of σ and α :

$$\lambda = 1.9, \quad \gamma = 1.4, \quad \eta_0 = 0.52, \quad \eta_{\pm 1} = 0.10.$$

Figure 10(a) shows in blue the value of the mean-squared error $\log \epsilon_{ms}(\sigma)$ as a function of $\log \sigma$ for $\alpha = 2$, computed on average over $\mathbf{C}^\alpha$ geometrically regular images. This error is below the error produced by the dyadic wavelet transform estimator, which shows the importance of the scattering $\mathbf{L}^1$ norms. More important, it has a different slope equal to $4/3$. It indicates that $\epsilon_{ms}(\sigma) \sim \sigma^{4/3}$ up to a $|\log \sigma|$ term that cannot be detected due to numerical uncertainties, which means that it has an asymptotic minimax rate. Figure 10(b) gives the value of $\log \epsilon_{ms}(\sigma)$ as a function of $\log \sigma$ for 4 Lipschitz exponents $\alpha \in \{2, 1.5, 1.2, 1\}$. We verify numerically that $\log \epsilon_{ms}(\sigma)$ has a linear increase in $\log \sigma$ which superimposes with lines of slope $2\alpha/(\alpha+1)$ shown in dashed line. It indicates that $\epsilon_{ms}(\sigma) \sim \sigma^{2\alpha/(\alpha+1)}$ up to a $|\log \sigma|$ term, which is the asymptotic minimax rate for all α . This shows that these scattering denoisers reaches the minimax rate for all $1 \leq \alpha \leq 2$, despite the presence of corners, and can thus adapt to unknown $\alpha \leq 2$.

The noise range on which the experiment is performed is identical for dyadic wavelet transform and scattering but differs from the one used for the UNet of Figure 3(b). Indeed, the UNet outperforms both estimators and is able to exploit the global shape of the deformed triangular foreground to denoise its contour.

Figure 11 illustrates the denoising results of dyadic wavelet transform and scattering transform for a particular image $f \in \Lambda$, for $\alpha = 2$. It is normalized between -1 and 1 and corrupted with a noise of variance $\sigma^2 = 0.1$. As expected, the mean-squared error improvement of the scattering estimator corresponds to the recovery of sharper edges, because the $\mathbf{L}^1$ norm of scattering coefficients give information on the edge profile. At corners, the scattering transform does not do any better than the wavelet estimators, but it has no impact on the asymptotic performances of the estimator because there is a finite number of corners which have no contribution to the asymptotic rate of the error. The UNet estimation shown in Figure 5 provides a much better recover of corners. A scattering transform estimator is much simpler but its performances are still far from the performances of a deep neural network estimator.

Conjecture on minimax scattering optimality Numerical denoising experiments indicate that a denoising estimator regularized by scattering $\mathbf{L}^1$ norms has a maximum error over $\mathbf{C}^\alpha$ geometrically regular images which reaches the asymptotic minimax rate. This is formalized by the following conjecture.

Conjecture 3.3. *Let Λ be a set of $\mathbf{C}^\alpha$ geometrically regular functions with bounded Lipschitz constants, for $1 \leq \alpha \leq 2$. Let $\hat{f}(g)$ be a scattering estimator of f contaminated by a Gaussian white noise of variance σ . There exists a constant C which does not depend upon α and σ such that*

$$\sup_{f \in \Lambda} \mathbb{E}_g(\|\hat{f}(g) - f\|^2) \leq C |\log \sigma| \sigma^{2\alpha/(\alpha+1)}. \quad (20)$$

The main difficulty to prove this conjecture is to prove the impact of the scattering $\mathbf{L}^1$ norms $\|\hat{f} * \psi_j^k\|_1$ for $k' \neq k^\perp$ on the denoising error. These $\mathbf{L}^1$ norms are maximised and not minimised as it is usually done in sparse representations, and it is the lines of constant phase of $\psi_j^{k'}$ which regularize $|\hat{f} * \psi_j^k|$ and hence the edge profile of $\hat{f}$. If the modulus is replaced by a rectifier, then the phase is replaced by the sign of real coefficients. Clearly, the




Figure 10: (a): Comparison between the mean-squared errors $\log \epsilon_{ms}(\sigma)$ of a dyadic wavelet estimator (red dashed curve) and a scattering estimator (blue dashed curve), computed on average over C^2 geometrically regular images defined in D.1. The blue and red dotted lines have a slope of 1 and 4/3. (b): mean-squared errors $\log \epsilon_{ms}(\sigma)$ for a scattering denoiser and C^α geometrically regular images, for different α . The different dashed coloured curve from red to blue corresponds to $\alpha \in \{2, 1.5, 1.2, 1\}$. These curves superimpose with the dotted lines having a slope $2\alpha/(\alpha + 1)$.




Figure 11: Top row: (a) uncorrupted C^2 -geometrically regular image discretised in dimension $d = 128^2$, ranging in $[-1, 1]$, (b) noisy version of this image, corrupted with an additive Gaussian white noise of variance $\sigma^2 = 0.1$, (c) denoised image obtained with a dyadic wavelet estimator, (d) denoised image obtained with a scattering estimator. Bottom row: images from the top row are zoomed on a region with a corner and edges ((e) original, (f) noisy, (g) dyadic wavelet estimator, (h) scattering estimator). The PSNR of the noisy image is improved from 16.01dB to 32.21 dB by the dyadic wavelet estimator and to 34.55 dB by the scattering estimator.

maximisation of these $\mathbf{L}^1$ norms provide a very different approach than minimisations of $\mathbf{L}^1$ norms of frame coefficients, to specify the geometric regularity of a function. The adaptivity of the resulting scattering estimator shows that it opens the possibility to create much more flexible estimators, as observed with deep neural networks. This conjecture gives a mathematical framework and a first step to better understand the mathematics underlying the performances of deep neural network estimators.

Acknowledgements This work was supported by PR[AI]RIE-PSAI-ANR-23-IACL-0008 and the DRUIDS projet ANR-24-EXMA-0002. It was granted access to the HPC resources of IDRIS under the allocations 2024-AD011014371R1 and 2025-AD011016159R1 made by GENCI. The authors thank Etienne Lempereur for his valuable feedback on the manuscript.

A Wavelet with directional vanishing moments

In the following we explain how to define a wavelet having directional vanishing moments in an angular cone. Lemma A.1 gives a sufficient condition to define the Fourier transform of such wavelets. We then detail how wavelets used in experiments of Section 3.3 are built.

Lemma A.1. *If the Fourier transform $\hat{\psi}$ of ψ is $\mathbf{C}^{m-1}$ and satisfies $\hat{\psi}(r_{-\theta}(\omega_1, \omega_2)) = O(|\omega_1|^m)$ for all $\theta \in [-\pi/4, \pi/4]$, where r_φ denotes the rotation of angle φ , then for any one-dimensional polynomial function $u_1 \mapsto q(u_1)$ of degree $m-1$:*

$$\forall \theta \in [-\pi/4, \pi/4], \quad \forall u_2 \in \mathbb{R}, \quad \int q(u_1) \psi(r_\theta(u_1, u_2)) du_1 = 0. \quad (21)$$

Moreover, ψ has m vanishing moments, i.e. for any two-dimensional polynomial function $(u_1, u_2) \mapsto q(u_1, u_2)$:

$$\int q(u_1, u_2) \psi(u_1, u_2) du_1 du_2 = 0.$$

Proof. Assume that $\hat{\psi}$ is $\mathbf{C}^{m-1}$. We first show that ψ has m directional vanishing moments in directions $\theta \in [-\pi/4, \pi/4]$. If $\hat{\psi}(\omega_1, \omega_2) = O(|\omega_1|^m)$ then $\partial_{\omega_1}^n \hat{\psi}(0, \omega_2) = 0$ for all $n \leq m-1$. It results that $\int u_1^n \psi(u_1, u_2) du_1 = 0$. Performing a rotation of angle θ proves that $\int u_1^n \psi(r_\theta(u_1, u_2)) du_1 = 0$ if $\hat{\psi}(r_{-\theta}(\omega_1, \omega_2)) = O(|\omega_1|^m)$, which is the first statement of the lemma. We prove that ψ has m vanishing moments using, again, that $\partial_{\omega_1}^n \hat{\psi}(0, \omega_2) = 0$ for all $n \leq m-1$. Indeed, it implies that all partial derivatives of $\hat{\psi}$ of order $n \leq m-1$ vanish at $(0, 0)$, from which we deduce the result. $\square$

Numerical experiments of Section 3.3 are performed with Morlet/Gabor wavelets [20, 43, 44], modified to fulfil conditions of Lemma A.1 with $m = 2$. The mother wavelet ψ is defined in the Fourier space as a product between the Fourier transform $\hat{\psi}_{\text{Morlet}}$ of a Morlet wavelet with two vanishing moments and a mask μ , defined in equation (23), satisfying $\mu(\omega_1, \omega_2) = 0$ for all (ω_1, ω_2) having a polar coordinate angle greater than $\pi/4$:

$$\hat{\psi} = \hat{\psi}_{\text{Morlet}} \times \mu. \quad (22)$$

The resulting wavelet ψ is $\mathbf{C}^\infty$, has two vanishing moments and two directional vanishing moments in the cone of angles $[-\pi/4, \pi/4]$. Figure 7 shows the real and imaginary parts of ψ , as well as its modulus and phase. Contrary to Theorem 3.2's hypothesis, ψ does not have a fast decay, $\hat{\psi}$ being only $\mathbf{C}^1$. This slowest asymptotic decay rate has no numerical impact, since we observe empirically that ψ inherits the spatial localization properties of ψ_{Morlet} for dimension $d = 128^2$ used in experiments of Section 3.3.

In what follows, we review the definition of Morlet wavelets and we define the mask μ . We then prove that the resulting wavelet ψ satisfies Lemma's A.1 hypothesis for $m = 2$.

Morlet wavelets A Morlet wavelet ψ_{Morlet} is defined as a product between a plane wave and a Gaussian envelop:

$$\psi_{\text{Morlet}}(u) = c_\sigma e^{-\|u\|^2/(2\sigma^2)} (e^{i\xi \cdot u} - k_{\sigma, \xi}),$$

where $x \cdot y$ denotes the standard scalar product in $\mathbb{R}^2$. The parameter $k_{\sigma, \xi}$ is a scalar, adjusted such that ψ has one vanishing moment: $\int \psi(u) du = 0$. Numerically, we use $\sigma = 0.7$ and $\xi = (1.05\pi, 0)$. A Morlet wavelet with two vanishing moments, still denoted ψ_{Morlet} , is defined by:

$$\psi_{\text{Morlet}}(u) = c_\sigma e^{-\|u\|^2/(2\sigma^2)} (e^{i\xi \cdot u} - k_{\sigma, \xi} - i\tilde{k}_{\sigma, \xi} \cdot u),$$

where $\tilde{k}_{\sigma,\xi}$ is a vector of $\mathbb{R}^2$ adjusted such that $\int u_1 \psi(u) du = 0$ and $\int u_2 \psi(u) du = 0$. Adjusting $\tilde{k}_{\sigma,\xi} = (\tilde{k}_{\sigma,\xi}^1, \tilde{k}_{\sigma,\xi}^2)$ is sufficient to impose vanishing partial derivatives of $\hat{\psi}_{\text{Morlet}}$ at $(0, 0)$ since for $i \in \{1, 2\}$:

$$\partial_{\omega_i} \hat{\psi}_{\text{Morlet}}(0, 0) = \sigma^2 \xi_i e^{-\sigma^2 \|\xi\|^2/2} - \frac{\sigma^2}{2\pi} \tilde{k}_{\sigma,\xi}^i.$$

The corresponding wavelet transform W , used in Section 3.3, satisfies a Littlewood-Paley formula for $\delta \in (0, 1)$, guaranteeing its invertibility:

$$\forall \omega \neq 0, 1 - \delta \leq |\hat{\phi}_{J_M}(\omega)|^2 + \sum_{j \leq J_M, k} |\hat{\psi}_j^k|^2(\omega) \leq 1 + \delta.$$

Constraint on the support of $\hat{\psi}$ The mask μ from equation (22) used in our numerical experiments is defined in polar coordinates as follows:

$$\forall r > 0, \mu(r, \varphi) = \tilde{\mu}(\varphi) \mathbb{1}_{|\varphi| < \pi/4} = \sin\left(\frac{\pi(1 + \sin(4\varphi + \pi/2))}{4}\right) \mathbb{1}_{|\varphi| < \pi/4} \text{ and } \forall \varphi \in [-\pi, \pi), \mu(0, \varphi) = 0. \quad (23)$$

It is $\mathbf{C}^1$ everywhere except at the origin and satisfies $\mu(r, \varphi) = 0$ for all $r > 0$ and $\varphi \notin [-\pi/4, \pi/4]$. Moreover, it has vanishing angular derivatives on the borders of the cone of angles $[-\pi/4, \pi/4]$:

$$\forall r > 0, \partial_\varphi \mu(r, \pi/4) = \partial_\varphi \mu(r, -\pi/4) = 0.$$

Denoting $(\mu_k)_k$ rotated versions of μ in directions $(k\pi/4)_{k < 4}$, this mask verifies: $\sum_k \mu_k^2 = 1$.

Fulfillment of Lemma's A.1 hypothesis To show that $\hat{\psi}$, defined through equation (22), satisfies Lemma's A.1 hypothesis for $m = 2$, we need to prove that it is $\mathbf{C}^1$ and that $\hat{\psi}(r_{-\theta}(\omega_1, \omega_2)) = O(\omega_1^2)$ for all $\theta \in [-\pi/4, \pi/4]$.

The Fourier transform $\hat{\psi}_{\text{Morlet}}$ and μ being respectively $\mathbf{C}^\infty$ on $\mathbb{R}^2$ and $\mathbf{C}^1$ on $\mathbb{R}^* \times \mathbb{R}^*$, $\hat{\psi}$ is $\mathbf{C}^1$ on $\mathbb{R}^* \times \mathbb{R}^*$. To prove that $\hat{\psi}$ is $\mathbf{C}^1$ on $\mathbb{R}^2$, one has to verify that its partial derivatives exist and are continuous in $(0, 0)$. Since $\hat{\psi}_{\text{Morlet}}(0, 0) = 0$ and since $\hat{\psi}_{\text{Morlet}}$ is a $\mathbf{C}^\infty$ function whose first order derivatives vanish at the origin, it satisfies $\hat{\psi}_{\text{Morlet}}(r, \varphi) = O(r^2)$ in polar coordinates. μ being bounded, this implies that $\hat{\psi}$'s partial derivatives exist and vanish at $(0, 0)$. $\partial_\varphi \mu$ being bounded too, it also implies that $\lim_{\omega \rightarrow (0,0)} \partial_{\omega_1} \hat{\psi}(\omega_1, \omega_2) = \lim_{\omega \rightarrow (0,0)} \partial_{\omega_2} \hat{\psi}(\omega_1, \omega_2) = 0$, which is proven expressing these partial derivatives in polar coordinates. $\hat{\psi}$ is thus $\mathbf{C}^1$ on $\mathbb{R}^2$.

Since μ 's support is included in the cone of angles $[-\pi/4, \pi/4]$, it is sufficient to show that $\hat{\psi}(r_{-\theta}(\omega_1, \omega_2)) = O(\omega_1^2)$ for $\theta = \pm\pi/4$ to get this result for all $\theta \in [-\pi/4, \pi/4]$. The indicator function $\mathbb{1}_{|\varphi| < \pi/4}$ does not play any role in the decay rate of $\hat{\psi}(r_{-\theta}(\omega_1, \omega_2))$ for $\omega_1 \rightarrow 0$. The result can thus be obtained studying the partial derivatives along ω_1 of the function $\hat{\psi}(\omega_1, \omega_2) \mapsto (\hat{\psi}_{\text{Morlet}} \times \tilde{\mu})(r_{-\theta}(\omega_1, \omega_2))$, where $\tilde{\mu}$ is defined in equation (23). Expressing these derivatives in polar coordinates and using the fact that $\tilde{\mu}$ has vanishing angular derivatives for $|\varphi| \rightarrow \pi/4$ and that $\hat{\psi}_{\text{Morlet}}(r, \varphi) = O(r^2)$, one shows that $\partial_{\omega_1} \hat{\psi}(\omega_1, \omega_2) = 0$ for all ω_2 and that $\partial_{\omega_1}^2 \hat{\psi}(\omega_1, \omega_2) \leq C$ for all $\omega_2 \neq 0$. We thus have $\hat{\psi}(r_{-\theta}(\omega_1, \omega_2)) = O(\omega_1^2)$ for all $\omega_2 \neq 0$. The case $\omega_2 = 0$ is given by $\hat{\psi}_{\text{Morlet}}$ decay for $\omega \rightarrow 0$.

B Proof of Theorem 3.1

Theorem. Suppose that ψ is a $\mathbf{C}^2$ function with two vanishing moments and two directional vanishing moments in the cone of angles $[-\pi/4, \pi/4]$, such that $|\psi|$ decays faster than any rational function. Let $\alpha \in [1, 2]$. Let f be a periodic image on $[0, 1]^2$ which is uniformly Lipschitz α above and below an edge curve, which is itself Lipschitz α and whose tangents have an angle in $[(k-1)\pi/4, (k+1)\pi/4]$. For all $s < \alpha$, there exists C such that:

$$\forall j \leq j_M, \forall u \in [0, 1]^2, |f * \psi_j^k(u)| \leq C 2^{(s-1)j}. \quad (24)$$

Proof. If $\alpha = 1$, the result is immediate, f being bounded and ψ_j^k being $\mathbf{L}^1$ normalized. In what follows we assume that $\alpha \in (1, 2]$. Thanks to ψ fast decay, for all $\epsilon > 0$, there is a constant $C > 0$ such that:

$$\int_{\|u\| > C_\psi 2^{(1-\epsilon)j}} |\psi_j(u)| du \leq C 2^{\alpha j}, \quad (25)$$

where C_ψ only depends on ψ . In the following, we consider $\epsilon \leq 1/2$. One has:

$$|f * \psi_j^k(u)| \leq \left| \int_{\|v\| \leq C_\psi 2^{(1-\epsilon)j}} f(u-v) \psi_j^k(v) dv \right| + \|f\|_\infty \int_{\|v\| > C_\psi 2^{(1-\epsilon)j}} |\psi_j^k(v)| dv. \quad (26)$$

The second integral is controlled thanks to equation (25), which shows that it is $O(2^{\alpha j})$. In the following, we prove that the first integral is $O(2^{(1-\epsilon)(\alpha-1)j})$. There are two cases depending whether the edge curve γ intersects or not the ball $B(u, C_\psi 2^{(1-\epsilon)j})$. If γ does not intersect $B(u, C_\psi 2^{(1-\epsilon)j})$, $f|_{B(u, C_\psi 2^{(1-\epsilon)j})}$ is uniformly $\mathbf{C}^\alpha$ and can be extended to a uniformly $\mathbf{C}^\alpha$ function $\tilde{f}$ on $[0, 1]^2$. Cutting the integral inside and outside $B(u, C_\psi 2^{(1-\epsilon)j})$, one gets:

$$\left| \int_{\|v\| \leq C_\psi 2^{(1-\epsilon)j}} f(u-v) \psi_j^k(v) dv \right| \leq \left| \int \tilde{f}(u-v) \psi_j^k(v) dv \right| + \|f - \tilde{f}\|_\infty \int_{\|v\| > C_\psi 2^{(1-\epsilon)j}} |\psi_j^k(v)| dv \leq C 2^{\alpha j}, \quad (27)$$

where the first integral is bounded using ψ vanishing moments through equation (7) and the second one with equation (25). Finally, combining equations (26) and (27):

$$|f * \psi_j^k(u)| \leq C 2^{\alpha j}.$$

If γ intersects $B(u, C_\psi 2^{(1-\epsilon)j})$, the hypothesis on γ 's tangent imposes that $B(u, C_\psi 2^{(1-\epsilon)j})$ contains only one connected edge portion, which can be parametrized along the direction $k\pi/4$. Without loss of generality, we assume a horizontal parametrization in what follows, i.e. $k = 0$. Let $\tilde{\gamma}$ be the Taylor expansion of order 1 of γ in $B(u, C_\psi 2^{(1-\epsilon)j})$:

$$\|\gamma - \tilde{\gamma}\|_\infty \leq C 2^{(1-\epsilon)\alpha j}.$$

The function $\tilde{\gamma}$ being linear, $(u_1, \tilde{\gamma}(u_1))$ is a straight line, whose angle with the horizontal axis is denoted θ . Thanks to ψ directional vanishing moments, one can consider without loss of generality the case $\theta = 0$. In this case, $|\gamma(\cdot)| = O(2^{(1-\epsilon)\alpha j})$ on $B(u, C_\psi 2^{(1-\epsilon)j})$. One has:

$$\left| \int_{\|v\| \leq C_\psi 2^{(1-\epsilon)j}} f(u-v) \psi_j^0(v) dv \right| \leq \int_{v_2 = -C_\psi 2^{(1-\epsilon)j}}^{C_\psi 2^{(1-\epsilon)j}} \left| \int_{v_1 = -C_\psi 2^{(1-\epsilon)j}}^{C_\psi 2^{(1-\epsilon)j}} f(u_1 - v_1, u_2 - v_2) \psi_j^0(v_1, v_2) dv_1 \right| dv_2. \quad (28)$$

Thanks to ψ_j^0 directional vanishing moments along the horizontal axis, the integral on v_1 could be seen as a one-dimensional wavelet coefficient of $v_1 \mapsto f(u-v)$ for a fixed v_2 . This integral is not taken on the whole space, but one could proceed as in equation (27) to exploit the vanishing moments. As illustrated in Figure 12(a), $v_1 \mapsto f(u-v)$ is discontinuous for $v_2 \in [-C 2^{(1-\epsilon)\alpha j}, C 2^{(1-\epsilon)\alpha j}]$, and $\mathbf{C}^\alpha$ otherwise. The corresponding one-dimensional wavelet coefficients would be $O(1)$ and $O(2^{\alpha j})$ respectively, for an L1 normalized wavelet. ψ_j^0 being two dimensional, one should count an additional 2^{-j} normalisation factor:

$$\left| \int_{v_1 = -C_\psi 2^{(1-\epsilon)j}}^{C_\psi 2^{(1-\epsilon)j}} f(u_1 - v_1, u_2 - v_2) \psi_j^0(v_1, v_2) dv_1 \right| \leq \begin{cases} C (2^{\alpha j} + 2^{-j}) & \text{if } v_2 \in [-C 2^{(1-\epsilon)\alpha j}, C 2^{(1-\epsilon)\alpha j}], \\ C (2^{\alpha j} + 2^{(\alpha-1)j}) & \text{otherwise,} \end{cases}$$

where the $2^{\alpha j}$ term accounts for $v_1 \notin B(0, C_\psi 2^{(1-\epsilon)j})$, using equation (25). Integrating on v_2 in equation (28) thus leads to:

$$\left| \int_{\|v\| \leq C_\psi 2^{(1-\epsilon)j}} f(u-v) \psi_j^0(v) dv \right| \leq C \left[2^{(1-\epsilon)j} 2^{(\alpha-1)j} + 2^{(1-\epsilon)\alpha j} \cdot 2^{-j} \right] \leq C 2^{[(1-\epsilon)\alpha-1]j}.$$

Injecting this bound in equation (26) and replacing $(1-\epsilon)\alpha - 1, \epsilon > 0$, by $s - 1, s < \alpha$, gives equation (13). $\square$

C Proof of Theorem 3.2

Theorem 3.2 states that the decay rate of scattering coefficients of the form $\| |f * \psi_j^k| * \psi_j^{k^\perp} \|_1$ depends on the regularity of edge curves. It relies on two main ideas:

- first, as shown by Theorem 3.1, the $\mathbf{C}^\alpha$ regularity of an edge can be captured by a wavelet exhibiting enough directional vanishing moments along the edge tangent, if $\alpha \in [1, 2]$. Lemma A.1 shows how to build wavelets with vanishing moments in a cone of directions, imposing a constraint on their Fourier transform support;




Figure 12: (a): Representation of f on $B(u, C_\psi 2^{(1-\epsilon)j})$, for u a point on the edge curve. The tangent of γ at u is assumed to be horizontal. Bounding $|f * \psi_j^0|(u)$ requires to exploit (1) the C^α regularity of $f_{u,v_2} : v_1 \mapsto f(u-v)$ for v_2 such that f_{u,v_2} does not intersect γ , (2) the fact that f_{u,v_2} is discontinuous only for a fraction $O(2^{(1-\epsilon)\alpha j})$ of v_2 possible values (between the two dashed lines). (b): Representation of $|f * \psi_j|$ on $B(u, C_\psi 2^{(1-\epsilon)j'})$. The tangent of γ at u is assumed to be vertical. The gray area corresponds to $O(1)$ wavelet coefficients $|f * \psi_j|(v)$ located along the edge. This region has a width $O(2^{(1-\epsilon)j})$ and is $O(2^{(1-\epsilon)j'})$ long. Bounding $|\rho(f * \psi_j) * \psi_{j'}^\perp|(u)$ requires to exploit (1) $\rho(f * \psi_j)$ small amplitude far from the edge (outside the region coloured in gray), (2) the fact that $v_2 \mapsto f(u-v)$ is Lipschitz 1 with a constant controlled thanks to γ regularity for a fraction $O(2^{(1-\epsilon)\alpha j'} + 2^{(1-\epsilon)j})$ of v_1 possible values (between the two dashed lines).

- second, since a wavelet cannot have vanishing moments in all directions, the ℓ^1 norm $\|\cdot * \psi^k\|_1$ will not capture edges regularity due to edge tangents ‘misaligned’ with ψ^k vanishing moments. One can then apply a second wavelet $\psi^{k^\perp}$ with complementary directional vanishing moments to get rid of the constraint of tangent alignment.

For the sake of clarity, the theorem is proved for the mother wavelet ψ , which oscillates along the horizontal axis. This is done without loss of generality, since one can always get back from ψ^k to ψ applying a rotation of angle $-k\pi/4$. The rotation of angle $\pi/2$ of ψ is denoted $\psi^\perp$. Dilated versions of ψ and $\psi^\perp$ at each scale 2^j are denoted ψ_j and $\psi_j^\perp$ respectively. They are L1 normalized:

$$\forall j \leq j_M, \|\psi_j\|_1 = 1.$$

The unit square $[0, 1]^2$ is mapped to the torus and periodized wavelets [40] are considered. Unless explicitly stated otherwise, the Euclidean distance is used. In what follows, C is a generic constant which may change from line to line. Theorem 3.2’s proof is in three stages, which are summarized hereafter in a proof sketch.

Directional vanishing moments and edges’ regularity Let f be a $C^\alpha[0, 1]^2$ geometrically regular function, $\alpha \in [1, 2]$, and let $u \in [0, 1]^2$. If u is far from f ’s edges, the wavelet coefficient modulus $|f * \psi_j|(u)$ is controlled thanks to ψ vanishing moments, as exposed in Section 2.2. If u is in the vicinity of an edge, one needs to exploit ψ directional vanishing moments. Suppose that the edge tangent at u forms an angle θ with the horizontal axis and that ψ has more than 2 vanishing moments along θ . Then, Theorem 3.1 proves that the corresponding edge coefficient $|f * \psi_j|(u)$ have a $O(2^{(1-\epsilon)(\alpha-1)j})$ decay rate, $\epsilon > 0$, instead of the $O(1)$ obtained without directional vanishing moments. For ψ having vanishing moments in the cone of directions $[-\pi/4, \pi/4]$, edge coefficients are split in two categories:

- those for which the edge tangent is in the cone, which have an amplitude $O(2^{(1-\epsilon)(\alpha-1)j})$;
- the others, which have an amplitude $O(1)$.

Edge coefficients are thus partially sparsified. One should still treat ‘misaligned’ edges.

Second wavelet transform Misaligned edge coefficients correspond to edges whose tangents are in the complementary of the cone of angle $[-\pi/4, \pi/4]$. The convolution $f * \psi$ inherits the geometric regularity of f , but it cannot be captured directly using the $\psi^\perp$'s directional vanishing moments. Indeed, such wavelets verify

$$\forall j \leq j' \leq j_M, \psi_{j'}^\perp * \psi_j = 0.$$

A non-linearity ρ should be added, leading to coefficients of the form $|\rho(f * \psi_j) * \psi_{j'}^\perp|$, where ρ must be chosen in order to ensure that $\rho(f * \psi_j)$ is still directionally regular. Lemma C.1 proves that, for ρ being Lipschitz 1 and verifying $\rho(0) = 0$, the scaling of such second order coefficients is in $O(2^{(1-\epsilon)(\alpha-1)jj'})$ in the neighbourhood of edges. In particular, the result holds for $\rho(\cdot) = |\cdot|$ and $\rho(\cdot) = \text{ReLU}(\cdot)$.

Theorem's proof For given $j \leq j' \leq j_M$ and $\epsilon > 0$, three types of coefficients should be considered:

- coefficients scaling in $O(2^{\alpha j})$ in uniformly regular regions, which are of area $O(1)$;
- coefficients scaling in $O(2^{(1-\epsilon)(\alpha-1)j})$ in neighbourhoods of edges, which are of area $O(2^{(1-\epsilon)j})$;
- coefficients scaling in $O(1)$ in neighbourhoods of junctions, which are of area $O(2^{2(1-\epsilon)j})$.

Thus, for all $\tilde{\epsilon} > 0$, there exists $C > 0$ such that $\|\rho(f * \psi_j) * \psi_{j'}^\perp\|_1 = C2^{(1-\tilde{\epsilon})\alpha j}$. Replacing $(1 - \tilde{\epsilon})\alpha$, $\tilde{\epsilon} > 0$, with $s < \alpha$ and summing on j, j' gives equation (15).

C.1 Decay of second order scattering coefficients

Lemma C.1. (partly inspired by Peyré and Mallat [45], proposition 1) Suppose that ψ satisfies the same hypothesis as in Theorem 3.1. Let f be a $\mathbf{C}^\alpha$ geometrically regular image with $\alpha \in [1, 2]$ with Lipschitz constant L and ρ be a Lipschitz 1 function verifying $\rho(0) = 0$. Let denote $\mathcal{C}_\delta \subset [0, 1]^2$ the set of points being δ -close from f 's corners. For all $s < \alpha$ and $\delta > 0$, there exists C such that:

$$\forall j \leq j' \leq j_M, \forall u \in [0, 1]^2 \setminus \mathcal{C}_{C_\psi 2^{(1-\delta)j'}}, |\rho(f * \psi_j) * \psi_{j'}^\perp|(u) \leq C2^{(s-1)j'}, \quad (29)$$

where C_ψ is defined in equation (25).

Proof. If $\alpha = 1$, the result is immediate, $\rho(f * \psi_j)$ being bounded and $\psi_{j'}^\perp$ being $\mathbf{L}^1$ normalized. In what follows, we thus consider $\alpha \in (1, 2]$. Let $u \in [0, 1]^2$ and $j \leq j' \leq j_M$. For all $\epsilon > 0$, one has:

$$|\rho(f * \psi_j) * \psi_{j'}^\perp|(u) \leq \left| \int_{\|v\| \leq C_\psi 2^{(1-\epsilon)j'}} \rho(f * \psi_j)(u - v) \psi_{j'}^\perp(v) dv \right| + \|\rho(f * \psi_j)\|_\infty \int_{\|v\| > C_\psi 2^{(1-\epsilon)j'}} |\psi_{j'}^\perp(v)| dv. \quad (30)$$

As in the proof of Theorem 3.1, the second integral is controlled thanks to equation (25). Using that $\rho(f * \psi_j)$ is bounded, one gets:

$$\|\rho(f * \psi_j)\|_\infty \int_{\|v\| > C_\psi 2^{(1-\epsilon)j'}} |\psi_{j'}^\perp(v)| dv = O(2^{\alpha j'}). \quad (31)$$

If f is uniformly $\mathbf{C}^\alpha$ over the ball $B(u, C_\psi [2^{(1-\epsilon)j} + 2^{(1-\epsilon)j'}])$, then, for all $v \in B(0, C_\psi 2^{(1-\epsilon)j'})$, it is uniformly $\mathbf{C}^\alpha$ over $B(u - v, C_\psi 2^{(1-\epsilon)j})$. Proceeding as in equation (27) thus leads to:

$$\forall v \in B(0, C_\psi 2^{(1-\epsilon)j'}), |f * \psi_j|(u - v) \leq C2^{\alpha j}. \quad (32)$$

The function ρ being Lipschitz 1 with $\rho(0) = 0$ and $\psi_{j'}^\perp$ being $\mathbf{L}^1$ normalized, we deduce from equation (32) a bound on the first integral of equation (30):

$$\left| \int_{\|v\| \leq C_\psi 2^{(1-\epsilon)j'}} \rho(f * \psi_j)(u - v) \psi_{j'}^\perp(v) dv \right| \leq C2^{\alpha j}. \quad (33)$$

Let now assume that f is not uniformly $\mathbf{C}^\alpha$ over $B(u, C_\psi [2^{(1-\epsilon)j} + 2^{(1-\epsilon)j'}])$. The Lipschitz constant L of f is finite and, by definition, its edge curves do not intersect tangentially. Moreover, u is $C_\psi 2^{(1-\delta)j'}$ -distant from f 's corners and all of the above holds for any $\epsilon > 0$. Then, for j, j' small enough and even if it means increasing a multiplicative

constant independent of the scale, the ball $B\left(u, C_\psi \left[2^{(1-\epsilon)j} + 2^{(1-\epsilon)j'}\right]\right)$ contains at most one connected edge portion from an edge curve γ , which could be parametrized either horizontally or vertically. Without loss of generality, we assume a horizontal parametrization. The angle of its best linear approximation with the horizontal axis is denoted θ . There are two cases, depending whether θ belongs to $[-\pi/4, \pi/4]$ or not. If $\theta \in [-\pi/4, \pi/4]$, one can apply Theorem 3.1 to get:

$$\forall s < \alpha, \exists C > 0 / \forall v \in B\left(0, C_\psi 2^{(1-\epsilon)j'}\right), |f * \psi_j|(u - v) = C 2^{(s-1)j}. \quad (34)$$

Finally, using the same arguments as for equation (33), we deduce from equation (34) that:

$$\left| \int_{\|v\| \leq C_\psi 2^{(1-\epsilon)j'}} \rho(f * \psi_j)(u - v) \psi_{j'}^\perp(v) dv \right| \leq C 2^{(s-1)j} \leq C 2^{(s-1)j'}. \quad (35)$$

If $\theta \notin [-\pi/4, \pi/4]$, the underlying geometric regularity will be captured by $\psi_{j'}^\perp$ instead of ψ_j . Thanks to ψ directional vanishing moments, we assume without loss of generality that $\theta = \pi/2$, meaning that $\psi_{j'}^\perp$ is aligned with the direction of regularity. Since ρ is not assumed to be differentiable, $f_j := \rho(f * \psi_j)$ does not inherit fully the regularity of f and ψ_j . Nevertheless, it inherits from $f * \psi_j$ a Lipschitz constant along the edge tangent. This constant is small, in a sense that will be specified, for j' not too big compared to j . The idea of the proof is to exploit this regularity for $j' < j/\alpha$ and to use the fact that f_j is of order 1 only on a small fraction of the support of $\psi_{j'}^\perp$ for $j' \geq j/\alpha$. One could proceed as in equation (27) to prove the existence of a constant $C > 0$ such that, for all $v \in B\left(0, C_\psi 2^{(1-\epsilon)j'}\right)$, $f_j(u - v) \leq C 2^{\alpha j}$ if the edge curve γ does not intersect $B(u - v, C_\psi 2^{(1-\epsilon)j})$. Denoting Γ the graph of the edge curve γ and using that $\|\psi_{j'}^\perp\|_1 = 1$, one gets:

$$\left| \int_{\|v\| \leq C_\psi 2^{(1-\epsilon)j'}} f_j(u - v) \psi_{j'}^\perp(v) dv \right| \leq \left| \int_{\substack{\|v\| \leq C_\psi 2^{(1-\epsilon)j'} \\ B(u-v, C_\psi 2^{(1-\epsilon)j}) \cap \Gamma \neq \emptyset}} f_j(u - v) \psi_{j'}^\perp(v) dv \right| + C 2^{\alpha j}. \quad (36)$$

Figure 12(b) illustrates in gray the set $u - \left\{v \in [0, 1]^2 / \|v\| \leq C_\psi 2^{(1-\epsilon)j'} \text{ and } B(u - v, C_\psi 2^{(1-\epsilon)j}) \cap \Gamma \neq \emptyset\right\}$. One has:

$$\int_{\substack{\|v\| \leq C_\psi 2^{(1-\epsilon)j'} \\ B(u-v, C_\psi 2^{(1-\epsilon)j}) \cap \Gamma \neq \emptyset}} dv = O(2^{(1-\epsilon)j + (1-\epsilon)j'}). \quad (37)$$

The function f_j is bounded and $\|\psi_{j'}^\perp\|_\infty = O(2^{-2j'})$, since it is $\mathbf{L}^1$ normalized. Using equation (37) and assuming that $j' \geq j/\alpha$, one thus gets:

$$\left| \int_{\substack{\|v\| \leq C_\psi 2^{(1-\epsilon)j'} \\ B(u-v, C_\psi 2^{(1-\epsilon)j}) \cap \Gamma \neq \emptyset}} f_j(u - v) \psi_{j'}^\perp(v) dv \right| \leq C 2^{(1-\epsilon)(j-j') - 2\epsilon j'} \leq C 2^{[(1-\epsilon)(\alpha-1) - 2\epsilon]j'}. \quad (38)$$

Replacing $(1-\epsilon)(\alpha-1) - 2\epsilon$, $\epsilon > 0$, by $s-1$, $s < \alpha$, one gets:

$$\forall s < \alpha, \exists C > 0 / \left| \int_{\|v\| \leq C_\psi 2^{(1-\epsilon)j'}} f_j(u - v) \psi_{j'}^\perp(v) dv \right| \leq C 2^{(s-1)j'}. \quad (39)$$

This bound does not exploit the directional vanishing moments of ψ . It is not suitable for $j' < j/\alpha$, since then $j - j'$ will no longer be upper bounded by $(\alpha-1)j'$. For $j' < j/\alpha$, one needs to exploit the directional regularity of f_j to control the remaining integral in equation (36) bound. The first order approximation $\tilde{\gamma}$ of γ have been assumed to be oriented vertically ($\theta = \pi/2$). As in the proof of Theorem 3.1, the maximal distance between γ and $\tilde{\gamma}$ over $B(u, C_\psi 2^{(1-\epsilon)j'})$ is controlled thanks to γ 's C^α regularity:

$$\|\gamma - \tilde{\gamma}\|_\infty \leq C 2^{(1-\epsilon)\alpha j'}.$$

Thus, there exists an interval $I_{j,j'} := \left[-C \left(2^{(1-\epsilon)j} + 2^{(1-\epsilon)\alpha j'}\right), C \left(2^{(1-\epsilon)j} + 2^{(1-\epsilon)\alpha j'}\right)\right] \subset \left[-C_\psi 2^{(1-\epsilon)j'}, C_\psi 2^{(1-\epsilon)j'}\right]$ such that:

$$\left\{v \in [0, 1]^2 / \|v\| \leq C_\psi 2^{(1-\epsilon)j'} \text{ and } B(u - v, C_\psi 2^{(1-\epsilon)j}) \cap \Gamma \neq \emptyset\right\} \subset I_{j,j'} \times \left[-C_\psi 2^{(1-\epsilon)j'}, C_\psi 2^{(1-\epsilon)j'}\right]. \quad (40)$$

The region $I_{j,j'} \times [-C_\psi 2^{(1-\epsilon)j'}, C_\psi 2^{(1-\epsilon)j'}]$ is localized between two vertical dashed lines in Figure 12(b). From equation (40), one gets:

$$\left| \int_{B(u-v, C_\psi 2^{(1-\epsilon)j}) \cap \Gamma \neq \emptyset} f_j(u-v) \psi_{j'}^\perp(v) dv \right| \leq \int_{v_1 \in I_{j,j'}} \left| \int_{v_2 = -C_\psi 2^{(1-\epsilon)j'}}^{C_\psi 2^{(1-\epsilon)j'}} f_j(u-v) \psi_{j'}^\perp(v) dv_2 \right| dv_1 + C 2^{\alpha j}, \quad (41)$$

with f_j being directionally regular along v_2 (since $\theta = \pi/2$). The $C 2^{\alpha j}$ term accounts for points $v \in I_{j,j'} \times [-C_\psi 2^{(1-\epsilon)j'}, C_\psi 2^{(1-\epsilon)j'}]$ such that $B(u-v, C_\psi 2^{(1-\epsilon)j}) \cap \Gamma = \emptyset$, for which $f_j(u-v) \leq C 2^{\alpha j}$. As in equation (28) from Theorem 3.1's proof, the integral on v_2 could be seen as a 1D wavelet coefficient of $v_2 \mapsto f_j(u-v)$ for a fixed v_1 , thanks to ψ directional vanishing moments. This integral is not taken on the whole space, but one could proceed as in equation (27) to exploit $\psi_{j'}^\perp$ directional vanishing moments and control the integral for $v_2 \notin [-C_\psi 2^{(1-\epsilon)j'}, C_\psi 2^{(1-\epsilon)j'}]$. We will use the Lipschitz regularity of this 1D function to get a bound on its wavelet coefficients, as illustrated in Figure 12(b). Its Lipschitz constant, which depends on v_1 , will be controlled through a bound on $\left| \frac{\partial f * \psi_j}{\partial u_2} \right|$. The function f is differentiable since $\alpha > 1$. Using triangular and Taylor-Lagrange inequalities, we get:

$$|f_j(u_1 - v_1, u_2 - v_2) - f_j(u_1 - v_1, u_2)| \leq C_{j,j'}(v_1) C_\rho |v_2|,$$

where $C_{j,j'}(v_1) := \left| \max_{v_2 \in [-C_\psi 2^{(1-\epsilon)j'}, C_\psi 2^{(1-\epsilon)j'}]} \frac{\partial f * \psi_j(u_1 - v_1, u_2 - v_2)}{\partial v_2} \right|$ and where C_ρ is the Lipschitz constant of ρ .

The function $v_2 \mapsto f_j(u-v)$ for a fixed v_1 is thus Lipschitz 1 with Lipschitz constant $C_{j,j'}(v_1) C_\rho$ for $v_2 \in [-C_\psi 2^{(1-\epsilon)j'}, C_\psi 2^{(1-\epsilon)j'}]$. Thanks to $\psi_{j'}^\perp$ directional vanishing moments along v_2 , we get:

$$\left| \int_{v_2 = -C_\psi 2^{2j'}}^{C_\psi 2^{2j'}} f_j(u-v) \psi_{j'}^\perp(v) dv_2 \right| \leq C \times C_{j,j'}(v_1) + C 2^{\alpha j'},$$

where the $2^{\alpha j'}$ term accounts for $v_2 \notin [-C_\psi 2^{(1-\epsilon)j'}, C_\psi 2^{(1-\epsilon)j'}]$, using equation (25). For an L1-normalized 1D wavelet, the bound would be proportional to $2^{j'}$. Here, this $2^{j'}$ factor, linked to f_j Lipschitz 1 regularity, compensates for the $2^{-2j'}$ normalization of $\psi_{j'}^\perp$. Combining this bound with equation (41) and integrating over v_1 leads to:

$$\left| \int_{B(u-v, C_\psi 2^{(1-\epsilon)j}) \cap \Gamma \neq \emptyset} f_j(u-v) \psi_{j'}^\perp(v) dv \right| \leq C \times C_{j,j'}^{\text{edge}} (2^{(1-\epsilon)j} + 2^{(1-\epsilon)\alpha j'}), \quad (42)$$

where $C_{j,j'}^{\text{edge}} = \max_{v_1 \in I_{j,j'}} C_{j,j'}(v_1)$. One thus needs to control $C_{j,j'}^{\text{edge}}$, which accounts to bound the derivative of $\tilde{u} \mapsto f * \psi_j(\tilde{u})$ along the edge tangent. The computation is similar to the proof of proposition 1 from Peyré and Mallat [45], the bound on $f * \psi_j$ derivatives being here obtained on a square region of length $2^{j'}$, $j' < j/\alpha$, instead of $2^{j/\alpha}$. Since $\alpha > 1$, γ and f along γ are differentiable, and one has:

$$\frac{\partial f * \psi_j}{\partial \tilde{u}_2}(\tilde{u}) = \frac{\partial}{\partial \tilde{u}_2} \int f(\tilde{u}_1 - v_1, \tilde{u}_2 - v_2) \psi_j(v_1, v_2) dv_1 dv_2 = \int \frac{\partial}{\partial \tilde{u}_2} [f(S(\tilde{u}, v)) \psi_j(T(\tilde{u}, v))] dv_1 dv_2,$$

where $S(\tilde{u}, v) := (\tilde{u}_1 + \gamma(\tilde{u}_2 - v_2) - v_1, \tilde{u}_2 - v_2)$ and $T(\tilde{u}, v) := (v_1 - \gamma(\tilde{u}_2 - v_2), v_2)$. The function f being regular along γ , $\left| \frac{\partial f(S(\tilde{u}, v))}{\partial \tilde{u}_2} \right|$ is uniformly bounded for $v_2 \in [-C_\psi 2^{(1-\epsilon)j}, C_\psi 2^{(1-\epsilon)j}]$ and $\tilde{u}_2 \in [-C_\psi 2^{(1-\epsilon)j'}, C_\psi 2^{(1-\epsilon)j'}]$.

Moreover, γ being C^α , one has for all $\tilde{u}_2 \in [-C_\psi 2^{(1-\epsilon)j'}, C_\psi 2^{(1-\epsilon)j'}]$:

$$\forall v_2 \in [-C_\psi 2^{(1-\epsilon)j}, C_\psi 2^{(1-\epsilon)j}], \left| \frac{\partial [\psi_j(T(\tilde{u}, v))]}{\partial \tilde{u}_2} \right| = \left| \gamma'(\tilde{u}_2 - v_2) \frac{\partial \psi_j}{\partial \tilde{u}_1}(T(\tilde{u}, v)) \right| \leq C 2^{(1-\epsilon)(\alpha-1)j'} 2^{-3j}.$$

By construction, f is bounded and $\|\psi_j\|_\infty = 2^{-2j}$. Integrating on v and using ψ_j 's fast decay through equation (25) then leads to:

$$\left| \frac{\partial f * \psi_j}{\partial \tilde{u}_2}(\tilde{u}) \right| \leq C \int_{\|v\| \leq C_\psi 2^{(1-\epsilon)j}} \left(|\psi_j(T(\tilde{u}, v))| + \left| \frac{\partial \psi_j(T(\tilde{u}, v))}{\partial \tilde{u}_2} \right| \right) dv + C 2^{\alpha j} \leq C 2^{-(1+2\epsilon)j + (1-\epsilon)(\alpha-1)j'}.$$

Injecting the result in equation (42) and using that $j' < j/\alpha$, one gets:

$$\left| \int_{\substack{\|v\| \leq C_\psi 2^{(1-\epsilon)j'} \\ B(u-v, C_\psi 2^{(1-\epsilon)j}) \cap \Gamma \neq \emptyset}} f_j(u-v) \psi_{j'}^\perp(v) dv \right| \leq C 2^{[(1-\epsilon)(\alpha-1)-3\epsilon\alpha]j'}. \quad (43)$$

Replacing $(1-\epsilon)(\alpha-1)-3\epsilon\alpha$, $\epsilon > 0$, by $s-1$, $s < \alpha$, one gets from equations (36) and (43):

$$\forall s < \alpha, \exists C > 0 / \left| \int_{\|v\| \leq C_\psi 2^{(1-\epsilon)j'}} f_j(u-v) \psi_{j'}^\perp(v) dv \right| \leq C 2^{(s-1)j'}. \quad (44)$$

Finally, combining equations (31), (33), (35), (39) and (44) with the bound of equation (30) gives the expected result. Equations (35), (39) and (44) have been established assuming that j, j' were small enough such that $B(u, C_\psi [2^{(1-\epsilon)j} + 2^{(1-\epsilon)j'}])$ contains at most one connected edge portion. This limiting scale depends only on f , which is fixed. Equation (29) thus remains true for j, j' above this scale, even if it means increasing the constant. $\square$

C.2 Main proof

Theorem. Suppose that ψ satisfies the same hypothesis as in Theorem 3.1. If f is a C^α geometrically regular function with $1 \leq \alpha \leq 2$ then, for all $s < \alpha$, there exists C such that:

$$\forall j \leq j' \leq j_M, \forall k < 4, \|f * \psi_j^k * \psi_{j'}^{k\perp}\|_1 \leq C 2^{sj'},$$

and

$$\forall s < \alpha, \forall k < 4, \sum_{j=-\infty}^{j_M} \sum_{j'=j+1}^{[j/\alpha]} 2^{-sj'} \|f * \psi_j^k * \psi_{j'}^{k\perp}\|_1 < +\infty.$$

Proof. Let $j \leq j' \leq j_M$ and $\epsilon > 0$. The unit square $[0, 1]^2$ is divided in three distinct regions $\mathcal{R}_{j,j'}^\epsilon$, $\mathcal{E}_{j,j'}^\epsilon$ and $\mathcal{C}_{j,j'}^\epsilon$, corresponding to varying degrees of regularity of f . The set $\mathcal{R}_{j,j'}^\epsilon$ contains points $u \in [0, 1]^2$ such that f is uniformly regular on the ball $B(u, C_\psi [2^{(1-\epsilon)j} + 2^{(1-\epsilon)j'}])$. Combining equations (30), (31) and (33) from Lemma C.1's proof, one obtains:

$$\exists C > 0 / \forall u \in \mathcal{R}_{j,j'}^\epsilon, \|f * \psi_j * \psi_{j'}^\perp(u)\| \leq C 2^{\alpha j'}. \quad (45)$$

The set $\mathcal{E}_{j,j'}^\epsilon$ contains points such that one and only one edge portion intersects $B(u, C_\psi [2^{(1-\epsilon)j} + 2^{(1-\epsilon)j'}])$. Proceeding as in Lemma C.1's proof and noticing that the leading term in the bound is given by equation (43), one gets:

$$\exists C > 0 / \forall u \in \mathcal{E}_{j,j'}^\epsilon, \|f * \psi_j * \psi_{j'}^\perp(u)\| \leq C 2^{[(1-\epsilon)(\alpha-1)-3\epsilon\alpha]j'}. \quad (46)$$

Finally, the set $\mathcal{C}_{j,j'}^\epsilon$ is defined as the complementary in $[0, 1]^2$ of $\mathcal{R}_{j,j'}^\epsilon$ and $\mathcal{E}_{j,j'}^\epsilon$. For j, j' small enough, it contains only points at the vicinity of f 's corners. The function f being bounded, there is a constant $C > 0$ such that:

$$\forall u \in \mathcal{C}_{j,j'}^\epsilon, \|f * \psi_j * \psi_{j'}^\perp(u)\| \leq C. \quad (47)$$

Combining equations (45), (46) and (47) one thus gets:

$$\exists C > 0 / \|f * \psi_j * \psi_{j'}^\perp\|_1 \leq C \left[2^{\alpha j'} \int_{\mathcal{R}_{j,j'}^\epsilon} du + 2^{[(1-\epsilon)(\alpha-1)-3\epsilon\alpha]j'} \int_{\mathcal{E}_{j,j'}^\epsilon} du + \int_{\mathcal{C}_{j,j'}^\epsilon} du \right]. \quad (48)$$

The integral $\int_{\mathcal{R}_{j,j'}^\epsilon} du$ is bounded by the area of the unit square. By construction, f has a finite number of corner points,

from what one deduces that $\int_{\mathcal{C}_{j,j'}^\epsilon} du \leq C 2^{2(1-\epsilon)j'}$. Finally, since the length of edge curves is fixed for a given f , one

has: $\int_{\mathcal{E}_{j,j'}^\epsilon} du \leq C 2^{(1-\epsilon)j'}$. Combining these bounds with equation (48) leads to:

$$\|f * \psi_j * \psi_{j'}^\perp\|_1 \leq C_{\text{regular}} 2^{\alpha j'} + C_{\text{edge}} 2^{[(1-\epsilon)(\alpha-1)-3\epsilon\alpha]j' + (1-\epsilon)j'} + C_{\text{corners}} 2^{2(1-\epsilon)j'} \leq C 2^{[(1-\epsilon)\alpha-3\epsilon\alpha]j'}.$$

Replacing $(1 - \epsilon)\alpha - 3\epsilon\alpha$, $\epsilon > 0$, by $s < \alpha$ gives the first statement of the theorem, even if it means increasing the constant to cover j, j' for which $C_{j,j'}^\epsilon$ is not reduced to a union of balls centred on f 's corners:

$$\forall s < \alpha, \exists C > 0 / \forall j \leq j' \leq j_M, \|f * \psi_j\| * \psi_{j'}^\perp\|_1 \leq C 2^{sj'}. \quad (49)$$

In particular, for any $s' < \alpha$, choosing $s > s'$ in equation (49) gives a constant $C > 0$ such that:

$$\sum_{j=-\infty}^{j_M} \sum_{j'=j+1}^{\lfloor j/\alpha \rfloor} 2^{-s'j'} \|f * \psi_j\| * \psi_{j'}^\perp\|_1 \leq C \sum_{j=-\infty}^0 2^{(s-s')j/\alpha} < +\infty,$$

which proves the second statement of the theorem. □

D Numerics

D.1 Discretising geometry

Data generation Numerical experiments are performed on data generated following an adapted version of Kadkhodaie et al. [28]. It defines a random process whose realizations are geometrically $C^\alpha[0, 1]^2$ functions, projected on a finite dimensional space $\mathbf{V}$ of dimension d (fixed to 128×128 in our experiments). Samples live on a uniform 2D grid of $\sqrt{d}$ pixels width - the grid dimension, 2, and the space dimension, d , should not be confused. A realization $P_{\mathbf{V}}f$ of this process is obtained in four steps. The projections on the 2D grid of two uniformly regular $C^\alpha[0, 1]^2$ functions are first sampled to form the background and the foreground, respectively denoted B and F . The contours of the foreground shape, which are the projections of three $C^\alpha[0, 1]$ functions, are then sampled on a 1D grid of length $\sqrt{d}$. They are combined to create a binary mask M defining the shape of the foreground on the 2D grid. The foreground is eventually multiplied by the mask M and added to the background:

$$P_{\mathbf{V}}f = B + M \times F.$$

Denoting $\gamma_1, \gamma_2, \gamma_3$ the contours sampled on the 1D grid and indexing the 2D grid with $(n_1, n_2) \in \llbracket 1, \sqrt{d} \rrbracket^2$, the mask M is defined as follows:

$$M(n_1, n_2) = \prod_i r_{\theta_i} \left(\mathbb{1}_{n_2/\sqrt{d} \leq \gamma_1(n_1)} \right),$$

where r_φ is the rotation of angle φ . For each $i = 1, 2, 3$:

$$\theta_i \sim \mathcal{U}([(1 + 3(i - 1))2\pi/9, (2 + 3(i - 1))2\pi/9]).$$

This constraint on contours orientation prevents boundary issues, since most of the time the resulting edges will intersect to form a triangular shape. If this shape crosses the boundary of the grid anyway, this boundary is considered as a fourth contour. This additional contour being a straight line, it is C^∞ and has no impact on the asymptotic denoising performances. The projection on a uniform 2D grid of width $\sqrt{d}$ (resp. 1D grid of length $\sqrt{d}$) of a random uniformly C^α function is obtained convolving a sample of Gaussian white noise of dimension d (resp. $\sqrt{d}$) and a filter defined in the Fourier space as $(c + \omega_1^2 + \omega_2^2)^{-(\alpha+1)/2}$ (resp. $(c + \omega)^{-(\alpha+1)}$). The parameter c controls the Lipschitz constant of the sampled uniformly regular signals. These Lipschitz constants define the set $\Lambda \subset \mathbf{L}^2[0, 1]^2$ from which functions f are drawn. The previous sampling procedure defines a random process on the projection of such functions in dimension d . It induces average Lipschitz constants and contours length, which are reported in table 1 for the different values of α used in our numerical experiments. The concentration of total contours length around its average value justifies the usage of ϵ_{mms} as a proxy of ϵ_m . Figure 13 gives examples of sampled images for α varying in $\{1, 1.2, 1.5, 2\}$.

	$\alpha = 1$	$\alpha = 1.2$	$\alpha = 1.5$	$\alpha = 2$
Lipschitz constant of contours	1.81/0.432	3.38/0.853	4.10/1.07	11.4/2.83
Lipschitz constant of uniform regions	16.8/2.84	39.8/8.68	153/34.8	120/32.4
Contours length	1.91/0.282	1.84/0.261	1.81/0.264	1.79/0.260

Table 1: Lipschitz constant and total contours length of the sampled functions (mean/standard deviation), estimated over 1000 i.i.d. realizations in dimension $d = 128^2$.

Taken as one-dimensional objects, contours are regular $C^\alpha[0, 1]$ functions. Nevertheless, this regularity is lost when they are discretised on the two-dimensional grid, since stair case discontinuities appear. The central aspect of this work being to study the ability of the scattering estimator to exploit and to be adaptive to frontiers' regularity, this issue must be solved cautiously. To mitigate it, d pixels images are generated as the low frequency components of a two levels FWT, with Symlet4 wavelets, of $4^2 d$ pixels images.




Figure 13: Examples of $\mathbf{C}^\alpha[0, 1]^2$ images discretised in dimension $d = 128^2$ and normalized in $[-1, 1]$, for α varying in $\{1, 1.2, 1.5, 2\}$.




Figure 14: Scaling of $\|f * \psi_j^k * \psi_j^{k+1}\|_1$ for $j \in \{j_m, \dots, j_m + 7\}$ and f being a $\mathbf{C}^2$ geometrically regular image with constant background and foreground - in order to focus on the regularity of the contour - discretised in dimension $d = 1024^2$. Theorem 3.2 predicts a scaling with a slope of 2 for such coefficients, which is represented by the gray dotted line. Geometry being discretised, the finest scales do not behave as expected. The y-axis is log-scaled.

Noise range Discretization implies that geometric regularity properties of $\mathbf{C}^\alpha$ images used in Section 3.3 are verified only on a given range of scales. This scale range implies a noise range on which one could expect to observe the theoretical MSE slope of $2\alpha/(\alpha + 1)$. Indeed, observing such a slope requires the denoiser to have an uncertainty on the contour's location at a scale for which the contour is still $\mathbf{C}^\alpha$. For the smallest scales, the contour is no longer regular, as illustrated in Figure 14. Consequently, one must be cautious on the value of the gap that separates the background and the foreground. Indeed, if it is too high, the contour could be estimated directly, for instance through a thresholding. There will be no uncertainty on its location (due to the limited precision allowed by the discretization scale), except for very high noise levels, and the MSE will vary as if the image was uniformly regular². Conversely, if this gap is too low, the image will look uniformly regular up to a certain scale. If this scale is too low compared to the scale range on which the theoretical behaviour is expected, then the experimental MSE slope will not be correct. Correctly bounding the gap between background and foreground is thus necessary so that there exists a noise range over which to observe the correct MSE slope. The precise value of this range will then depend on the performance of the used denoiser. The better the denoiser, the more accurately it will be able to estimate the position of the contour at high noise levels, and

²For a wavelet thresholding denoising, all edge coefficients would be very high valued. They will thus get out quickly from the noise, letting only coefficients from uniform regions for smaller noise levels. This issue would not exist in a continuous setting, since there will be edge coefficients of all amplitudes.

the faster the uncertainty about its location will disappear. Thus, applied to the same data as the scattering estimator, the UNet whose performances are reported in Figure 3(b) has a $2\alpha/(\alpha + 1)$ MSE slope for noise levels much higher than those used for the Figure 10. In summary, given the finite size of the denoised images, the theoretical behaviour can only be observed over a range of upper and lower bounded noise. For C^α geometrically regular images, discretised in dimension $d = 128^2$, the gap value between background and foreground has been experimentally set between 0.4 et 0.6 for data normalized between -1 and 1 . These values were chosen ensuring that the dyadic wavelet estimator MSE has a slope of 1, while the scattering estimator gives the theoretically expected slope.

D.2 Denoisers and optimization

UNet architecture The UNet used for experiments of Figures 3(b) and 5 follows the architecture proposed in Ho et al. [24]. Implementation was obtained from <https://huggingface.co/blog/annotated-diffusion>. It has been trained on a dataset of 44000 $C^\alpha[0, 1]^2$ geometrically regular images discretised in dimension $d = 128^2$, composed of 11000 samples from each degree of regularity $\alpha \in \{1, 1.2, 1.5, 2\}$. It has been split into train (36000 samples), validation (4000 samples) and test (4000 samples) datasets. Early stopping with patience 20 and delta 0.1 was used. The network was trained with the Adam optimizer [29] and a MSE loss function.

Wavelet thresholding denoiser Figure 2(d) presents the result of a denoising by thresholding in a Symlet4 wavelet basis. Slopes from Figure 3(a) are obtained with a thresholding denoiser, again in a Symlet4 basis, average over 10 translations along each direction. In both cases, thresholding was performed following the rule prescribed by Donoho and Johnstone [14].

Scattering estimator Optimizing parameters $\{\lambda, \gamma, \eta_0, \eta_{\pm 1}\}$ in order to minimize the MSE of the variational denoiser of equation (2) could be written as a bi-level optimization problem [59]. More precisely, since we want to minimize its MSE for a whole range of noise, we consider the following problem:

$$\underset{\lambda, \gamma, \eta_0, \eta_{\pm 1}}{\operatorname{argmin}} \sum_{\sigma} \log(\mathbb{E}_{P_{\mathbf{V}} f, g}(\|\hat{f}_{\lambda, \gamma, \eta_0, \eta_{\pm 1}}(g) - P_{\mathbf{V}} f\|_2^2)), \text{ s.t. } \hat{f}_{\lambda, \gamma, \eta_0, \eta_{\pm 1}}(g) = \underset{h \in \mathbf{V}}{\operatorname{argmin}} \frac{1}{2} \|h - P_{\mathbf{V}} g\|^2 + \sigma^2 U_{\lambda, \gamma, \eta_0, \eta_{\pm 1}}(g). \quad (50)$$

The need to adapt the scale parameters j_m et j_M depending on the noise level σ prevents the usage of implicit differential methods to solve straightforwardly problem 50 [60, 17]. As the number of parameters to be optimized was limited, a grid search approach was adopted. For the dyadic wavelet estimator penalization term defined in equation (18), there is only one parameter to optimize: λ . This is done computing the MSE obtained for each λ of the grid, at each noise level $\sigma^2 \in \{1.05, 0.67, 0.43, 0.27, 0.18, 0.11, 0.07, 0.05\}$, for the best choice of j_M . The one maximising the criterion of equation (50) is then selected, leading to $\lambda = 1.2$. For the scattering penalization term defined in equation (19), the same procedure is applied to optimize $\{\lambda, \gamma, \eta_0, \eta_{\pm 1}\}$. The criteria of equation (50) is computed with three noise levels $\sigma^2 \in \{0.43, 0.11, 0.03\}$. One should notice a subtlety compared to the first order case: j_m requires also to be adjusted for the scattering estimator. Indeed, as the image size is not reduced according to the noise level, j_m must be adapted for the terms controlling the singularities' support, otherwise artificially sharp contours - badly positioned - would be reconstructed. Finally, the finer scale being special due to the discretization issues discussed in D.1, the corresponding scattering coefficients - those for which the first convolution is performed with the finest wavelets $\psi_{j_m}^k$ - are regularized with a multiplicative factor 0.55, which was adjusted empirically.

D.2.1 Variational denoising optimization

The optimization problem defined in equation (2) with U given by equation (19) is non-standard. Unlike energies defined in equations (4) or (18), the regularizer is non-convex and its proximal operator is unknown. This precludes the use of standard proximal methods such as proximal gradient descent [46], ISTA [9], and pFISTA [37]. The literature proposes extensions of these approaches for non-convex regularizers, but they require assumptions that are not verified by the scattering energy, such as the property to be written as a difference of convex functions [21]. Furthermore, while the scattering transform can be interpreted as a two layers neural network, Plug-and-Play (PnP) optimization methods [48] are not applicable in this context. PnP approaches typically employ neural networks to approximate MSE denoisers rather than energy functionals, with the denoiser subsequently integrated into an ADMM framework [50, 52] for denoising [33]. The neural network is generally not differentiated. Even in cases where neural networks are used to model energy functionals directly, as in Hurault et al. [25], modifications are required to ensure Lipschitz continuity of the gradient, which is not guaranteed for the scattering transform. We therefore evaluated several gradient-based optimization methods (L-BFGS [35], Stochastic Gradient Descent, and Adam [29]), ultimately selecting L-BFGS based on its superior performance. To ensure a fair comparison with the dyadic wavelet estimator defined in equation (18), we applied the same optimization algorithm to both approaches. Given the non-convex nature of the optimization

problem, we tested multiple initialization strategies: random initialization, initialization from the noisy image, and initialization from the image denoised using the dyadic wavelet estimator. These different initialization schemes yielded no significant differences in denoising performance.

References

- [1] Anestis Antoniadis and Jianqing Fan. Regularization of wavelet approximations. *Journal of the American Statistical Association*, 96(455):939–967, 2001.
- [2] Joan Bruna and Stephane Mallat. Invariant scattering convolution networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1872–1886, 2013. doi: 10.1109/TPAMI.2012.230.
- [3] Joan Bruna and Stephane Mallat. Multiscale sparse microcanonical models, 2019. URL <https://arxiv.org/abs/1801.02013>.
- [4] Emmanuel J. Candès and David L. Donoho. Curvelets: A surprisingly effective nonadaptive representation for objects with edges. In *Curve and Surface Fitting*, volume 2, pages 105–120, 1999.
- [5] Emmanuel J. Candès, Laurent Demanet, David L. Donoho, and Lexing Ying. Fast discrete curvelet transforms. *Multiscale Model. Simul.*, 5:861–899, 2006. URL <https://api.semanticscholar.org/CorpusID:7951288>.
- [6] Emmanuel Candès and David Donoho. New tight frames of curvelets and optimal representations of objects with c2 singularities. *Communications on Pure and Applied Mathematics*, 57:219 – 266, 02 2004. doi: 10.1002/cpa.10116.
- [7] Ronald R Coifman and David L Donoho. Translation-invariant de-noising. In *Wavelets and statistics*, pages 125–150. Springer, 1995.
- [8] Ingrid Daubechies. Ten lectures on wavelets. *Computers in Physics*, 6:697–697, 1992. URL <https://api.semanticscholar.org/CorpusID:58524360>.
- [9] Ingrid Daubechies, Michel Defrise, and Christine De Mol. An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Communications on Pure and Applied Mathematics*, 57, 2003. URL <https://api.semanticscholar.org/CorpusID:1438417>.
- [10] Minh N. Do and Martin Vetterli. Contourlets: a directional multiresolution image representation. *Proceedings. International Conference on Image Processing*, 1:I–I, 2002. URL <https://api.semanticscholar.org/CorpusID:12960781>.
- [11] David L Donoho. Unconditional bases are optimal bases for data compression and for statistical estimation. *Applied and computational harmonic analysis*, 1(1):100–115, 1993.
- [12] David L. Donoho. Wedgelets: nearly minimax estimation of edges. *Annals of Statistics*, 27:859–897, 1999. URL <https://api.semanticscholar.org/CorpusID:15209351>.
- [13] David L Donoho. De-noising by soft-thresholding. *IEEE transactions on information theory*, 41(3):613–627, 2002.
- [14] David L. Donoho and Iain M. Johnstone. Ideal spatial adaptation by wavelet shrinkage. *Biometrika*, 81:425–455, 1994. URL <https://api.semanticscholar.org/CorpusID:239520>.
- [15] David L Donoho, Iain M Johnstone, Gérard Kerkycharian, and Dominique Picard. Density estimation by wavelet thresholding. *The Annals of statistics*, pages 508–539, 1996.
- [16] David L. Donoho, Iain M. Johnstone, Jeffrey C. Hoch, and Alan S. Stern. Maximum entropy and the nearly black object. *Journal of the Royal Statistical Society: Series B (Methodological)*, 54(1):41–67, 12 2018. ISSN 0035-9246. doi: 10.1111/j.2517-6161.1992.tb01864.x. URL <https://doi.org/10.1111/j.2517-6161.1992.tb01864.x>.
- [17] Asen L Dontchev and R Tyrrell Rockafellar. *Implicit functions and solution mappings*, volume 543. Springer, 2009.
- [18] Charles Dossal, Erwan Le Pennec, and Stéphane Mallat. Bandlet image estimation with model selection. *Signal Process.*, 91:2743–2753, 2008. URL <https://api.semanticscholar.org/CorpusID:9079834>.
- [19] Michael Elad, Bahjat Kwar, and Gregory Vaksman. Image denoising: The deep learning revolution and beyond—a survey paper. *SIAM Journal on Imaging Sciences*, 16(3):1594–1654, 2023.
- [20] Dennis Gabor. Theory of communication. part 1: The analysis of information. *Journal of the Institution of Electrical Engineers-part III: radio and communication engineering*, 93(26):429–441, 1946.
- [21] Pinghua Gong, Changshui Zhang, Zhaosong Lu, Jianhua Huang, and Jieping Ye. A general iterative shrinkage and thresholding algorithm for non-convex regularized optimization problems. In *international conference on machine learning*, pages 37–45. PMLR, 2013.

- [22] Karol Gregor and Yann LeCun. Learning fast approximations of sparse coding. In *International Conference on Machine Learning*, 2010. URL <https://api.semanticscholar.org/CorpusID:13333501>.
- [23] Kanghui Guo, Gitta Kutyniok, and Demetrio Labate. Sparse multidimensional representations using anisotropic dilation and shear operators. *Wavelets and splines*, 14:189–201, 2006.
- [24] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [25] Samuel Hurault, Arthur Leclaire, and Nicolas Papadakis. Gradient step denoiser for convergent plug-and-play. *ArXiv*, abs/2110.03220, 2021. URL <https://api.semanticscholar.org/CorpusID:238419652>.
- [26] Stéphane Jaffard. Pointwise smoothness, two-microlocalization and wavelet coefficients. *Publicacions Matemàtiques*, 35:155–168, 1991. URL <https://api.semanticscholar.org/CorpusID:120072795>.
- [27] Zahra Kadkhodaie and Eero P Simoncelli. Stochastic solutions for linear inverse problems using the prior implicit in a denoiser. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=x5hh6N9bUUb>.
- [28] Zahra Kadkhodaie, Florentin Guth, Eero P Simoncelli, and Stéphane Mallat. Generalization in diffusion models arises from geometry-adaptive harmonic representations. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=ANvmVS2Yr0>.
- [29] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014. URL <https://api.semanticscholar.org/CorpusID:6628106>.
- [30] Aleksandr Petrovich Korostelev and Alexandre B Tsybakov. *Minimax theory of image reconstruction*. Lecture Notes in Statistics. Springer New York, NY, 1993.
- [31] Matthieu Kowalski. Thresholding rules and iterative shrinkage/thresholding algorithm: A convergence study. In *2014 IEEE International Conference on Image Processing (ICIP)*, pages 4151–4155. IEEE, 2014.
- [32] Demetrio Labate, Wang-Q Lim, Gitta Kutyniok, and Guido L. Weiss. Sparse multidimensional representation using shearlets. In *SPIE Optics + Photonics*, 2005. URL <https://api.semanticscholar.org/CorpusID:8829697>.
- [33] Rémi Laumont, Valentin De Bortoli, Andrés Almansa, Julie Delon, Alain Durmus, and Marcelo Pereyra. On maximum a posteriori estimation with plug & play priors and stochastic gradient descent. *Journal of Mathematical Imaging and Vision*, 65:140–163, 2022. URL <https://api.semanticscholar.org/CorpusID:246016443>.
- [34] E. Le Pennec and S. Mallat. Sparse geometric image representations with bandelets. *IEEE Transactions on Image Processing*, 14(4):423–438, 2005. doi: 10.1109/TIP.2005.843753.
- [35] Dong C. Liu and Jorge Nocedal. On the limited memory bfgs method for large scale optimization. *Mathematical Programming*, 45:503–528, 1989. URL <https://api.semanticscholar.org/CorpusID:5681609>.
- [36] Jialin Liu, Xiaohan Chen, Zhangyang Wang, and Wotao Yin. Alista: Analytic weights are as good as learned weights in lista. In *International Conference on Learning Representations*, 2018. URL <https://api.semanticscholar.org/CorpusID:53312645>.
- [37] Yunsong Liu, Zhifang Zhan, Jian-Feng Cai, Di Guo, Zhong Chen, and Xiaobo Qu. Projected iterative soft-thresholding algorithm for tight frames in compressed sensing magnetic resonance imaging. *IEEE transactions on medical imaging*, 35(9):2130–2140, 2016.
- [38] Shangzhen Luan, Chen Chen, Baochang Zhang, Jungong Han, and Jianzhuang Liu. Gabor convolutional networks. *IEEE Transactions on Image Processing*, 27(9):4357–4366, 2018.
- [39] Stéphane Mallat. A theory for multiresolution signal decomposition: The wavelet representation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 11:674–693, 1989. URL <https://api.semanticscholar.org/CorpusID:2356353>.
- [40] Stéphane Mallat. *A wavelet tour of signal processing*. Academic Press, 1999.
- [41] Stéphane Mallat. Group invariant scattering. *Communications on Pure and Applied Mathematics*, 65(10):1331–1398, 2012.
- [42] Yves Meyer. *Wavelets and operators*. Number 37 in Cambridge Studies in Advanced Mathematics. Cambridge university press, 1992.
- [43] Jean Morlet, Georges Arens, Eliane Fourgeau, and Dominique Glard. Wave propagation and sampling theory—part i: Complex signal and scattering in multilayered media. *Geophysics*, 47(2):203–221, 1982.
- [44] Jean Morlet, Georges Arens, Eliane Fourgeau, and Dominique Glard. Wave propagation and sampling theory—part ii: Sampling theory and complex waves. *Geophysics*, 47(2):222–236, 1982.

- [45] Gabriel Peyré and Stéphane Mallat. Orthogonal bandelet bases for geometric images approximation. *Communications on Pure and Applied Mathematics*, 61, 2008. URL <https://api.semanticscholar.org/CorpusID:8946083>.
- [46] Nicholas G Polson, James G Scott, and Brandon T Willard. Proximal algorithms in statistics and machine learning. *Statistical Science*, pages 559–581, 2015.
- [47] Zhuwei Qin, Fuxun Yu, Chenchen Liu, and Xiang Chen. How convolutional neural network see the world-a survey of convolutional neural network visualization methods. *arXiv preprint arXiv:1804.11191*, 2018.
- [48] Yaniv Romano, Michael Elad, and Peyman Milanfar. The little engine that could: Regularization by denoising (red). *SIAM Journal on Imaging Sciences*, 10(4):1804–1844, 2017.
- [49] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. *ArXiv*, abs/1505.04597, 2015. URL <https://api.semanticscholar.org/CorpusID:3719281>.
- [50] Shirin Shoushtari, Weijie Gan, and Ulugbek S Kamilov. Convergence of nonconvex pnp-admm with mmse denoisers. In *2023 IEEE 9th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, pages 511–515. IEEE, 2023.
- [51] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=PxtTIG12RRHS>.
- [52] Yu Sun, Zihui Wu, Xiaojian Xu, Brendt Wohlberg, and Ulugbek S Kamilov. Scalable plug-and-play admm with convergence guarantees. *IEEE Transactions on Computational Imaging*, 7:849–863, 2021.
- [53] Rini Smita Thakur, Shubhojeet Chatterjee, Ram Narayan Yadav, and Lalita Gupta. Image de-noising with machine learning: A review. *IEEE Access*, 9:93338–93363, 2021.
- [54] Andrey N Tikhonov. Solution of incorrectly formulated problems and the regularization method. *Sov Dok*, 4: 1035–1038, 1963.
- [55] Andrey N. Tikhonov and Vasiliy Y. Arsenin. *Solutions of ill-posed problems*. V. H. Winston & Sons, Washington, D.C.: John Wiley & Sons, New York, 1977. Translated from the Russian, Preface by translation editor Fritz John, Scripta Series in Mathematics.
- [56] Norbert Wiener. *Extrapolation, Interpolation, and Smoothing of Stationary Time Series: With Engineering Applications*. The MIT Press, 08 1949. ISBN 9780262257190. doi: 10.7551/mitpress/2946.001.0001. URL <https://doi.org/10.7551/mitpress/2946.001.0001>.
- [57] Fisher Yu, Yinda Zhang, Shuran Song, Ari Seff, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015.
- [58] Matthew D. Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision – ECCV 2014*, pages 818–833, Cham, 2014. Springer International Publishing. ISBN 978-3-319-10590-1.
- [59] Yihua Zhang, Prashant Khanduri, Ioannis C. Tsaknakis, Yuguang Yao, Min-Fong Hong, and Sijia Liu. An introduction to bilevel optimization: Foundations and applications in signal processing and machine learning. *IEEE Signal Processing Magazine*, 41:38–59, 2023. URL <https://api.semanticscholar.org/CorpusID:260378880>.
- [60] Nicolas Zucchet and João Sacramento. Beyond backpropagation: Bilevel optimization through implicit differentiation and equilibrium propagation. *Neural Computation*, 34(12):2309–2346, November 2022. ISSN 1530-888X. doi: 10.1162/neco_a_01547. URL http://dx.doi.org/10.1162/neco_a_01547.